Technical Report 1149

Emotional Synthetic Forces

Amy E. Henninger
Soar Technology, Inc.

Eric Chown
University of Maine

Randy Jones
Soar Technology, Inc.

July 2004

United States Army Research Institute
for the Behavioral and Social Sciences

U.S. Army Research Institute for the Behavioral and Social Sciences

A Directorate of the U.S. Army Human Resources Command

ZITA M. SIMUTIS
Director

Research accomplished under contract
for the Department of the Army

Soar Technology, Inc.

Technical Review by

Carl W. Lickteig, U.S. Army Research Institute
Paula Durlach, U.S. Army Research Institute

NOTICES

DISTRIBUTION: Primary distribution of this Technical Report has been made by ARI. Please address correspondence concerning distribution of reports to: U.S. Army Research Institute for the Behavioral and Social Sciences, Attn: DAPE-ARI-PO, 2511 Jefferson Davis Highway, Arlington, Virginia 22202-3926

FINAL DISPOSITION: This Technical Report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The findings in this Technical Report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

REPORT DOCUMENTATION PAGE

1. REPORT DATE (dd-mm-yy)
July 2004

2. REPORT TYPE
Final

3. DATES COVERED (from. . . to)
September 1999 to January 2003

4. TITLE AND SUBTITLE
Emotional Synthetic Forces

5a. CONTRACT OR GRANT NUMBER
DASW01-99-C-0037

5b. PROGRAM ELEMENT NUMBER
2Q262785

5c. PROJECT NUMBER
A790

5d. TASK NUMBER
233

5e. WORK UNIT NUMBER
C01

6. AUTHOR(S)
Amy E. Henninger (Soar Technology, Inc.), Eric Chown (University of Maine), and Randy Jones (Soar Technology, Inc.)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Soar Technology, Inc.
3600 Green Court, Suite 600
Ann Arbor, MI., 48105

8. PERFORMING ORGANIZATION REPORT NUMBER
DASW01-99-C-0037

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Research Institute for the Behavioral and Social Sciences
2511 Jefferson Davis Highway
Arlington, VA 22202-3926

10. MONITOR ACRONYM
ARI

11. MONITOR REPORT NUMBER
Technical Report 1149

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES Contracting Officer's Representative: Bob G. Witmer.

14. ABSTRACT *(Maximum 200 words)*:
Report developed under a STTR contract number DASW01-99-C-0037 entitled Emotional Synthetic Forces. The objective of this research was to make the decision-making process of complex agents less predictable and more realistic, by incorporating emotional factors that affect humans. To this end, researchers adopted an approach that promotes the emergence of behavior as a result of complex interactions between factors affecting emotions, integrated in a connectionist-style model, and factors affecting decision making, represented in a symbolic model.

This report explains the model sub-components, integration, and testing, which includes a description of the behaviors we used in the development of prototype, the design of experiments, a representative set of behavior patterns that emerged as a result of exercising the model over the design space, and recommendations for future work. Results of prototype suggest an effective means of increasing the variability of computer-generated forces (CGF) behavior in manner consistent with modern day emotional theories.

15. SUBJECT TERMS
Emotions model, synthetic forces, computer-generated forces, CGF, human performance models, human behavior models, symbolic models, decision-making

SECURITY CLASSIFICATION OF

16. REPORT
Unclassified

17. ABSTRACT
Unclassified

18. THIS PAGE
Unclassified

19. LIMITATION OF ABSTRACT
Unclassified

20. NUMBER OF PAGES

21. RESPONSIBLE PERSON
(Name and Telephone Number)
Bob G. Witmer
DSN 970-3995
Comm 407-384-3995

Technical Report 1149

Emotional Synthetic Forces

Amy E. Henninger
Soar Technology, Inc.

Eric Chown
University of Maine

Randy Jones
Soar Technology, Inc.

Simulator Systems Research Unit
Stephen L. Goldberg, Chief

U.S. Army Research Institute for the Behavioral and Social Sciences
2511 Jefferson Davis Highway, Arlington, Virginia 22202-3926

July 2004

Army Project Number
20262785A790

Personnel Performance and Training Technology

FOREWORD

Computer-Generated Forces (CGF) found in most current combat simulations act more like drones or robots than like humans. They are too predictable and lack individual human traits and distinct personality types. They almost always have perfect knowledge of the battlefield and do not experience human emotions such as anger, or fear. Unlike the foes that our Soldiers will face on the battlefield their decisions and actions are not influenced by their arousal levels or emotions. Thus current CGF do not provide a realistic opponent for training. The U.S. Army Research Institute for the Behavioral and Social Sciences (ARI) has sought to remedy this training deficiency by funding through the STTR process three different approaches to representing human personality and emotions.

This report documents an approach that integrates a connectionist model of emotions from Chown (1993) into a robust Soar cognitive architecture (Rosenbloom, Laird, & Newell (1993). The emotions model includes pleasure/pain and clarity/confusion subsystems that differentiate between positive and negative emotional states, leading to varying arousal levels. Arousal then affects the decision process which is implemented within a symbolic cognitive architecture. Static and dynamic tests of this model suggest that the addition of the emotions model increases the variability of CGF performance, making the CGF less predictable – more like humans and less like robots.

ARI's Simulator Systems Research Unit is interested in making training more realistic. One way to do this is by improving the human performance models used to control CGF behaviors. The work described is part of ARI research task 202a, Virtual Environments Research for Infantry Training and Simulation. This research has been briefed to members of the military modeling community including PEO-STRI, the RDE Command, and NAVAIR-TSD, and was presented as a paper at the 23rd Army Science Conference.

Stephen L. Goldberg
Acting Technical Director

EMOTIONAL SYNTHETIC FORCES

EXECUTIVE SUMMARY

Research Requirement:

The Army needs to train against realistic synthetic agents in simulated environments to prepare for the real combatants that they will encounter on the battlefield. The objective of this research is to make the decision-making process of complex agents less predictable and more realistic, by incorporating emotional factors that affect humans. The application area incorporates emotions and individual differences into the behavior models of synthetic virtual special-forces agents in a battlefield simulation. This is an ideal test area for a model of emotions, because the intelligent agents must exercise a variety of reasoning capabilities, including situation assessment, planning, reacting to goal failures, and interacting with a team of agents. While the framework for this model is being developed within the military domain, we anticipate that the design is general enough to apply to other areas (e.g., animated characters) as well.

Procedure:

In tune with modern theories of emotions (e.g., Damasio, 1995; LeDoux, 1992), we regard emotions essentially as subconscious signals and evaluations that inform, modify, and receive feedback from a variety of sources including higher cognitive processes and the sensorimotor system. Thus, our work explicitly distinguishes the subconscious processes (in a connectionist implementation) and the decision making that is subject to emotional influences (in a symbolic cognitive architecture). This was accomplished by integrating a connectionist model of emotions from Chown (1993) with Rosenbloom, Laird, and Newell's (1993; Newell, 1990) Soar architecture.

It is our position that "emotional states" are emergent patterns of interaction between decision-making knowledge and these emotional signal systems. To this end, we have adopted an approach that promotes the emergence of behavior as a result of complex interactions between factors affecting emotions, integrated in the connectionist-style model, and factors affecting decision making, represented in the symbolic model.

In our framework, symbolic assessments of a small set of "emotional attributes" reside in a working memory. These attributes include indicators of affect, arousal, and information degradation. Portions of working memory are accessible by the deliberate cognitive process, and portions are accessible by the emotion mechanisms, so memory serves as the interface between the two. These working memory elements combine with background knowledge to generate strategies, reasoning, and external behavior. At the same time, the cognitive model creates working interpretations of the environment and status of internal goals (situational awareness). Some of these interpretations and assessments feed into the connectionist model, which in turn continuously computes new values for each emotional attribute. Instead of considering cognition and emotion as opposing forces, this architecture supports the view that they evolved together to maintain effective responses to stimuli that influence the survival of the self and the species.

Findings:

Experiments and analyses demonstrate that the incorporation of an emotions model can make the behavior of CGF less predictable. Specifically, in the prototype detailed in this report, the use of an emotions model increased the variability in the agent's response an average of 3.1. Further, this increase in response space is non-probabilistic, so responses can be easily understood and interpreted by operators. Moreover, what makes this approach a particularly elegant way for generating less predictable behavior is the fact that the increase in response variability is a result of an additional input that is internal to the Agent and thus, not detectable by humans interacting with the scenario. Thus, while operators can easily debug the behavior of the agent, trainees cannot easily "game" the system.

Utilization of Findings:

Based on our research, we can conclude that the use of an emotions model can increase behavior variability in reasonable ways. Thus, the use of the emotions model could be added to simulations that rely on a believable and unpredictable CGF for training our Soldiers. However, we have not demonstrated that emotional IFOR behavior is more realistic, nor have we demonstrated that the use of emotional IFORs will improve training. It is our opinion and recommendation that somewhere along this vein of research, funding agencies and sponsors of behavior moderator research formally investigate the assumed benefits of incorporating these models into IFOR systems.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

AAA	Anti-aircraft Artillery
APC	Armored Personnel Carrier
ARI	Army Research Institute
CGF	Computer Generated Force
FM	Field Manual
FEBA	Forward Edge of Battle Area
FOB	Forward Operating Base
FSM	Finite State Machine
FWA	Forward Wing Aircraft
IFOR	Intelligent Forces
IBMs	Individual-based models
LRR	Long Range Reconnaissance
MAS	Multi-Agent System
ORP	Objective Rally Point
PEN	Psychoticism, Extraversion, Neuroticism
PIR	Priority Intelligence Requirements
RF	Radio Frequency
RP	Release Point
RWA	Rotary Wing Aircraft
SAF	Semi-Automated Forces
SATCOM	Satellite Communications
SME	Subject Matter Expert
SMI	Soar ModSAF Interface
SOF	Special Operations Forces
TRACE	Tracing Recurrent Activity in Cognitive Elements

Introduction

This paper describes the development, implementation, and evaluation of a framework for modeling emotions in an interactive, decision-making agent. The primary goal of this work is to make the decision-making process of complex agents more realistic with regard to behavior moderators, including emotional factors, which affect humans. That is, instead of being uniform agents that rely entirely on a deterministic body of expertise to make their decisions, the decision making processes of our agents vary according to select emotional factors affecting the agent, including the agent's parameterized emotional profile. In tune with modern theories of emotions (e.g., Damasio, 1995; LeDoux, 1992), we regard emotions essentially as subconscious signals and evaluations that inform, modify, and receive feedback from higher cognitive processes. Thus, our work explicitly distinguishes the subconscious processes (in a connectionist implementation) and the decision making that is subject to emotional influences (in a symbolic cognitive architecture). The premise of this model is that emotions serve as a kind of automatic assessment system that can guide or otherwise influence the more deliberative decision making process.

Our framework integrates a connectionist model of emotions from Chown (1993) with Rosenbloom, Laird, and Newell's (1993; Newell, 1990) Soar cognitive architecture. The primary components of this emotional system are pleasure/pain and clarity/confusion subsystems that differentiate between positive and negative states. These, in turn, feed into an arousal system that interfaces with the decision-making system. It is our position that "emotional states" are in reality emergent patterns of interaction between decision-making knowledge and these emotional signal systems. To this end, we have adopted an approach that promotes the emergence of behavior as a result of complex interactions between factors affecting emotions, integrated in the connectionist-style model, and factors affecting decision making, represented in the symbolic cognitive architecture.

The remainder of this report documents the development and evaluation of this approach. The section entitled Related Research provides context for the work. The Architecture section, explains how our model was derived and implemented. Then, next two main sections Experimentation and Discussion, respectively, present the methods used to evaluate our research. Finally, the Conclusions section, offers lessons learned and strategies for future development.

Related Research

Literature review for this project resulted in the review of a number of scientific papers and books. Reviews of these articles have been incorporated in past progress reports submitted to ARI. In the interest of time and space, only those articles considered essential to the context and development of our model are presented in this report. Specifically, in this report, we present like models developed by other researchers in the military simulation community, a very brief historical overview of emotions research in the psychological community, and review of current emotions research in cognitive science community.

Emotions in Military Simulations

A recent panel report sponsored by the National Research Council has called for the use of personality factors, behavior moderators and emotions to develop more realistic CGFs (Pew and Mavor, 1998). These recommendations have spawned a number of studies incorporating fatigue representations (French, 2001; Jones et al., 1997), defeat mechanisms (Heeringa & Cohen, 2000), personality paradigms (Hudlicka & Billingsley, 1999; McKenzie et al., 2001), and emotion models (Fransechini et al., 2001; Gratch & Marsella, 2001; Hudlicka & Billingsley, 1999) in prominent CGF systems. While we currently know of no studies that have investigated whether military training is improved by the use of CGFs with these capabilities, Army instructional

courseware designers have recognized the significance of emotions in learning and training (Abell, 2000).

Our work has some similarities to the work of Gratch and Marsella (2001). In large part, this is due to the fact that both systems make use of the Soar architecture for decision-making. Thus, they share some similarity in the constructs for relating an agent's emotions its decisions. However, there are also many important differences in the systems. For example, the model of emotional intensity presented in this paper is influenced by more factors. Gratch has discrete high-level emotions (e.g., fear and anger), where such high-level constructs emerge from patterns of behavior in our model. Also, our model can be influenced by individual differences in temperament. On the other hand, Gratch uses a more generic approach to emotional appraisal.

Hudlicka and Billingsley (1999) also make the connection between emotional content and leadership style by representing the effects of temperament in their framework. However, emphasis appears to be placed on the influence of emotions on decision-making, with little focus placed on the complexity or variability of the emotional intensity model.

Alternatively, Fransechini et al (2001) place a much greater emphasis on deriving emotional intensity through a highly biologically based, neurophysiological approach. The model is based on two dimensions including arousal and distress, but its precise form is not discernible from current publications.

Given the immaturity of the research in this area, we are not able to conclusively state that one approach is better than another, only that differences exist. But, we expect that the models developed from this generation of research will be useful in assessing whether military training is indeed improved through the integration of emotional models into synthetic forces.

Emotions Modeling History

The study of emotion has had a fickle history in psychology, which appears to be correlated to prominent psychological theories of the day (Schultz, 1981). Over the last two decades, rapid growth in our understanding of brain function and in how it relates to behavior has renewed interest in emotion as a research area. Also, exciting progress in experimental neurobiology paralleled by explosive development of connectionist models has contributed to the resurgence of emotions research. The term "connectionist", coined by psychologists, is used to convey the fact that many psychological constructs are better explained in terms of distributed, parallel networks of adaptive units as opposed to terms of serial symbolic processing units. Practically speaking, a connectionist system can be thought of as the application of neural networks to high-level cognition (Barnden, 1995). A variety of neural network studies have already begun to address a wide range of issues (e.g., motivation, emotion, and goal direction) in cognition and behavior (Levine, 1992). Interestingly, many of the concepts of connectionist psychology are strongly related to work in behaviorism, where the former provides a stronger "internal structure" using simple units with explicit learning rules rather than simple stimulus-response probabilities.

There is currently no universally accepted, comprehensive theory of emotions. Instead, there exist a host of "mini-theories" that emphasize cognitive, motivational, physiological, and behavioral dimensions of emotion. For example, cognitive theorists tend to focus on thoughts and evaluations when defining emotions, physiologists tend to focus on physiological reactions, behaviorists on emotional behavior, and so on. While individual camps exist, there is now a growing list of researchers (Lazarus, 1984; Ortony, 1988; Levine and Leven, 1992) who generally support the notion that emotional states can be manipulated by a combination of *different* factors. At a minimum, these factors seem to include cognitive processes (expectations) and physiological states (usually interpreted as arousal).

A second concept that is common to many emotion theories is the existence of a central evaluative mechanism that determines whether a given situation is potentially harmful or beneficial to the individual. For example, LeDoux and Fellous (1995) have discovered neural circuitry that processes stimuli according to whether they threaten or enhance the survival of the organism or its species. Also, a related discovery of an emotional memory system that works in concert with this circuitry has further added to the recent thrust of emotion research. Emotional memory has been associated with the amygdala and appears to add an "emotional flavor" to a declarative memory, which is thought to primarily originate in the hippocampus. This theory, exercised at its most primitive level, suggests that emotions are strong, "hard-wired" responses to stimuli that have a positive or negative survival value. The accompanying work on emotional memory suggests that these responses are mostly learned through classical conditioning (LeDoux, 1992) and performed as unconscious processes (Damasio, 1995).

Clearly, the models of emotions proposed in the psychological community are not only complex, but still in their formative stages. This gives rise to a system that is difficult to express in computational terms. However, there are some consistencies among the theories, and it is our strategy to use these generally accepted common themes to the extent possible. In those cases where no one theme prevails, we adopt a more functional, physiologically based approach as it tends to be more readily expressed in computational terms.

Model Adopted for Research

Cognitive science has a long tradition as viewing emotion as being disruptive to rational thought (e.g., Salovey & Mayer, 1990; Tversky, & Kahneman, 1974; Nisbett & Ross, 1980). However, emotions can also be viewed as an efficient mechanism to change the decision making process in beneficial ways. As Kaplan et al (1991) noted, human survival relies upon information processing rather than sharp claws or teeth, so humans have developed numerous ways to process information efficiently. Our model adopts the position that emotions are correlated to survival value. Thus, the model extends Kaplan et al.'s work by building on the premise that primitive emotional responses enhance survival and that more complex emotions (e.g., those based on cognition) should then serve the same purpose. In this instance, a primitive emotional response such as "fearing a bear" is treated the same way as a cognitive emotional response such as "fearing a gun". In both cases fear is an appropriate response, useful for avoiding potentially dangerous situations.

The most basic way that humans code information in emotional terms is by the valence of experiences, where some are coded as being positive (pleasurable) and others are coded as negative (painful). In general terms, pleasure positively correlates to survival of the self or of the species, while pain has a negative correlation to survival. In turn, the intensity of pleasure and pain serves to measure the strength of this correlation and indicates the relative importance of the experience. In our theory, the role of the emotional system is to respond appropriately to pleasure or pain, or to the prediction of pleasure or pain. With the addition of cognitive structure, mixtures of pleasure and pain, etc., this system gets much more complex, but the fundamental idea remains the same.

In humans, the primary mechanism for responding to important events is the arousal system. What is commonly called arousal is actually a collection of related responses including, among others, increased heart rate and respiration, and changes in levels of dopamine and other chemicals in the brain. As such, in our system, the basic function of the arousal system is to modulate cognitive responses (such as decisions, attention, etc.) according to the valence and intensity of the current situation. The idea is that the chemical changes that occur during arousal impact the way in which the brain processes information. Different combinations of arousal and pleasure or pain will influence the precise changes in how information is processed. For example, if a thought or activity leads to pleasure the combination of pleasure and arousal should lead to changes that reinforce and continue the behavior. When experiencing pain, on the other hand, the

opposite is true, it is important to stop whatever is causing the pain. The strength of the response is typically proportional to the strength of the sensation.

There are a number of ways that the cognitive system builds on these basic mechanisms. Through evolution, certain stimuli may become important, by being intrinsically arousing, even if they have no inherent valence (e.g., snakes). In addition, through learning people can come to anticipate pleasure and pain based upon past experience. For example, anticipation of a physical sensation that has been pleasurable in the past can itself be pleasurable.

One other cognitive refinement as noted by Kaplan (1991) indicates that organisms that rely upon information for survival can also benefit from knowing how good their available information and relevant experience is. Actions are riskier when taken either with poor information or a poor world model. With respect to cognition, in situations where people have good models and good information they feel a sense of clarity. Conversely, with poor world models or a lack of information, people experience confusion. In cognitive terms, these are both highly important events, as they provide strong signals about the quality of one's experience with regard to the current situation. Someone who is confused is unlikely to make good decisions. Someone who has clarity, on the other hand, should be able to use their experience effectively. These signals can serve to guide behavior – when people are confused they may desire to retreat to more familiar, presumably safer, environments.

The constructs of arousal, valence (pleasure/pain), and information quality (clarity/confusion) proposed by Kaplan form the basis of the emotions model developed for this research. In the following section, the overall architecture of the emotions model and the decision-making model is presented, with specific detail on each of the models being provided initially and more detail on the interactions and interfaces of the individual models being provided last.

Architecture

The work of Kaplan described in the previous section is embodied in an architecture called Tracing Recurrent Activity in Cognitive Elements (TRACE). This architecture forms the basis of our emotions model, which was in turn, integrated with a rule-based model based on the Soar architecture. In the following sections, the theory behind each architecture is presented, as is the implementation of our model within those architectural principles. Specifically, the first section presents the emotions model based on TRACE and the specific derivation of the model used in our research. Similarly, the next section presents the agent model based on Soar and describes the specifics of that model as applicable to this research. Finally a section describing the interfaces required between these two models and the way in which these two models will interact is included.

Connectionist Model Representing Emotion

Building on the work of Kaplan presented in the previous section, this model uses arousal to distinguish the intensity of emotions, and uses pleasure and pain to distinguish positive affect emotions such as joy from negative affect emotions such as sadness. One of the benefits of this approach is that it is not necessary to posit specific mechanisms for differing emotions, as has been done in several other synthetic emotional systems (e.g., Velasquez, 1997; 1998; Gadanho & Hallam, 1998). Rather, specific emotional outcomes will arise out of different interactions of the general-purpose mechanisms we are constructing. Pleasure, pain, clarity, confusion, and arousal are all multi-purpose mechanisms that translate a complex variety of signals into fairly simple cognitive imperatives. For example, experiencing pain is a call for action to do something to stop the pain. As demonstrated in Figure 1, in this framework, emotional states are combinations of a number of components including arousal, pleasure/pain, and temporal factors. For example, "fear" is associated with high levels of arousal stemming from the anticipation of pain. Changing each

emotional component will result in the interpretation of a different emotional label. For example, at lower levels of arousal "fear" becomes "anxiety". Alternatively, if the arousal trigger is a past event instead of an anticipated event, the emotional interpretation might change to "remorse" or "anger" depending upon attentional factors relating to the source of the pain.

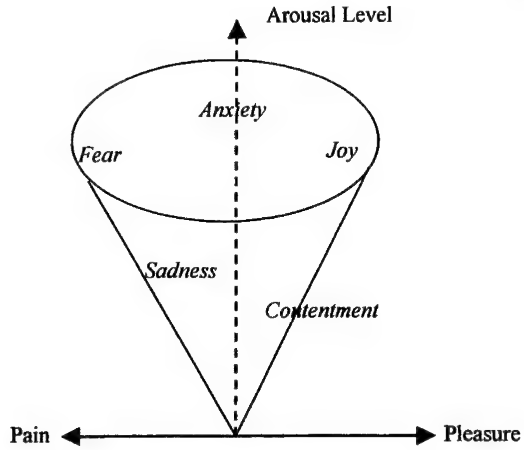

Figure 1. Interaction between arousal and valence.

<u>TRACE</u>

TRACE (Kaplan, Weaver, French, 1990) is a modern version of Hebb's cell assembly construct (Hebb, 1949; Kaplan, et al., 1991). A cell assembly is a collection of neurons that, due to their strong interconnections, tend to act as a unit. In this model, a given cell assembly acts as the internal representation for objects, whether concrete or abstract, in the world. When an object is perceived, the corresponding cell assembly becomes active. In this model, therefore, cell assemblies serve as the basic unit of thought. Learning occurs as the result of the continued activity of a cell assembly (reverberation).

Hebb developed the cell assembly construct to address questions concerning the temporal nature of neural processing. Essentially a cell assembly is a large collection of neurons which act in concert and which have temporal extent due to their recurrent connections and their corresponding ability to "reverberate." Hebb's theory lost favor initially in part because he omitted inhibition, a construct for which there was no evidence at the time. More recently, however, cell assemblies have undergone something of a revival as advances in neuroscience have been incorporated in the theory (Kaplan, et al., 1991) and experimental evidence for their existence has been found (Amit, 1995).

In TRACE, the emphasis is on simulating the internal dynamics of a population of neurons that would comprise a cell assembly. In TRACE various neural control mechanisms are postulated to play different functional roles in the cognitive system. For example, inhibition is useful as a selection mechanism when multiple cell assemblies are competing to become active. A major addition to cell assembly theory by the Kaplan group was to add fatigue to counterbalance the reverberation inherent in a highly recurrent system.

TRACE uses a set of difference equations that are updated at each time step to model the collective behavior of a large group of neurons. The equations model various biological functions such as activity, neural fatigue, short-term connection strength, long-term connection strength, sensitivity to firing, and network or external input. Each of these factors impacts the duration and intensity of neural firing with the cell assembly. In turn it is correlated neural firing that drives learning. One line of research on TRACE has been to show that a simple neural learning rule (the so called "Hebb rule") is capable of extremely sophisticated behavior when

placed in the context of a greater cognitive system, particularly the emotional system (Chown, 1994; Chown, 2002).

Kaplan et al. argued that units built with these basic properties have a number of advantages over the simple units used in many traditional connectionist models (1991). Different levels of activity in a cell assembly, for example, can serve different cognitive purposes, such as coding for conscious versus unconscious processing. The major questions left open by the original work on TRACE was how the notion of a single cell assembly could be extended to the cognitive system as a whole. This resulted in the expansion of TRACE into MultiTrace (Sonntag, 1991; Forbell & Chown, 2000; Chown, 2002). MultiTrace consists of multiple cell assemblies arranged in layered networks. Among other things, this development allows for sequence learning, a critical step in building cognitive structure.

<u>Emotions Model Derivation and Implementation</u>

Because we are postulating a few interacting mechanisms with continuous values, we have conceptualized this part of the architecture as a connectionist network. As described in the previous section, the connectionist model consists of several interacting components:

1) An arousal level system
2) A pleasure/pain system
3) A clarity/confusion system

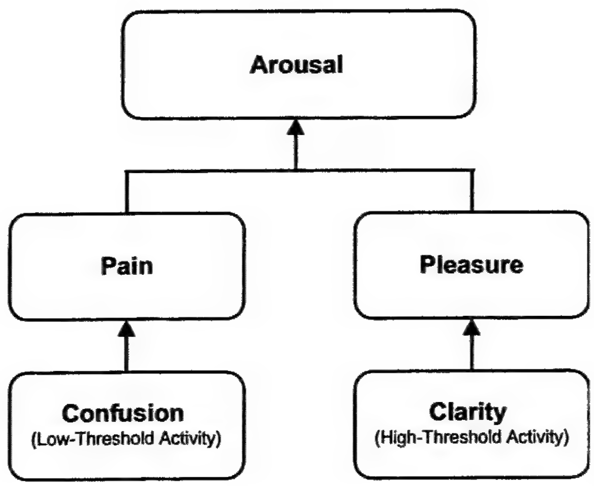

Figure 2. The emotional subsystem.

In our model, these mechanisms (pleasure/pain, arousal, and clarity/confusion) form the core of the emotional system. The advantage of such a generalized system is that it does not require specialized processing for every conceivable situation or emotion. In the next section we describe the theory in more detail, as well as an implementation of the emotional system.

<u>Clarity/Confusion</u>. Kaplan (1991) postulated the constructs of clarity and confusion as important correlates of pleasure and pain in forms of higher intelligence. Kaplan suggests that members of each species possess particular characteristics that facilitate their survival. For example, because humans are not particularly fast, fierce, or camouflaged, we rely on our ability to organize, store, and use information to enhance our survival. As a result, confusion is a potentially dangerous attribute and clarity is a desirable attribute, since an organism that is confused is less likely to respond in accord with its best interests. In this sense, clarity and confusion incorporate the influence of information quality and perceptual accuracy on the level of arousal. That is, the confusion/clarity mechanism becomes a measure that evaluates the relationship between the world and the person's knowledge of the world.

The derivation of the clarity/confusion components may be seen in equation 1.

$$
\begin{aligned}
&C = C_+ - C_- \\
&\text{where} \\
&C_+ = \eta_+ \sum \chi_{+,i} \\
&C_- = \eta_- \sum \chi_{-,i} \\
&\text{and where} \\
&\quad \chi_+ = \text{clarity inputs} \\
&\quad \chi_- = \text{confusion inputs} \\
&\quad \eta_+ = \text{susceptibility to pleasure constant} \\
&\quad \eta_- = \text{susceptibility to pain constant}
\end{aligned} \tag{1}
$$

Table 1.

Factors Influencing Clarity System

Clarity Inputs	**Value**
expectation-met **yes**	0.25
response-worked **yes**	0.4
received-required-information **yes**	0.2
leader-knows-I'm-alive **yes**	0.2
enemy-exists **no**	0.4
enemy-sees-me **no**	0.5

Table 2.

Factors Influencing Confusion System

Confusion Inputs	**Value**
Unknown value – benign	0.1
Unknown value – useful	0.25
Unknown value – dangerous	0.4
response-worked **no**	0.2
received-required-information **no**	0.2
input-overload **yes**	0.5
expectation-met **no**	0.1-1.0
leader-knows-I'm-alive **no**	0.4
enemy-exists **yes**	0.2
enemy-sees-me **yes**	0.2

Pain/Pleasure. The pleasure/pain continuum system is designed to interpret the level to which a stimulus represents a threat or enhancement to the survival of the species and consequently, will affect an organism's arousal level. Whereas undifferentiated arousal might yield undifferentiated activity, pleasure and pain provide additional information for an improved general response. In the case of pleasure the response should be an arrest response to continue the pleasurable stimuli – taking action might cause the stimuli to stop. For pain, on the other hand, the proper response is an excite response in order to do something to stop the painful stimuli.

Pleasure and pain help an organism make basic distinctions between beneficial and harmful stimuli. The adaptive benefits of a pleasure/pain system are virtually self-evident. At their most primitive levels painful stimuli are damaging to an organism while pleasurable stimuli are either replenishing or are oriented towards reproduction. Pleasure and pain, without the need for analysis, provide an organism with a strong message about its current state. Activities that bring pain need to be terminated quickly, while activities that bring pleasure should be continued. These signals are immediate and do not require any intermediate processing. By themselves pleasure and pain confer an adaptive advantage to organisms for the simple fact that pain should be avoided and pleasure extended. A creature that simply retracts a limb upon feeling pain has an adaptive advantage over one that must analyze the sensation and determine a rational course of action.

Pleasure and pain both denote important events, and therefore will increase an organism's arousal level. Whereas undifferentiated arousal might yield undifferentiated activity (for example agitation), pleasure and pain provide additional information for an improved general response. In the case of pleasure the response should be to *arrest* any new responses – taking a new action might cause the pleasure to stop. For pain, on the other hand, the proper response is to *excite* possible responses in order to do something to stop the painful stimulation. We note that these are tendencies that may be overridden by cognition or other factors. For example, in the face of overwhelming stimulation, an organism may simply freeze. In other cases, organisms may learn to respond in less natural ways.

By themselves, pleasure, pain and arousal form a useful, general purpose, system. At this point they have little or nothing to do with information processing and accordingly they serve all manner of organisms well, not just higher animals. For information processing organisms, such as humans, pleasure and pain can serve in a greatly expanded role. With cognition comes the ability to make predictions. Among the predictions an organism can make is whether or not it will experience pleasure or pain based upon its actions. As with Damasio's "somatic marker hypothesis" (1994), such predictions will be accompanied by bodily feelings commensurate with the results of the predictions. The ability to make predictions has been widely cited as the primary advantage of an information-processing organism. Interestingly enough, whereas the direct sensation of pleasure stimulates arrest, and the direct sensation of pain stimulates excitement, for predicted pleasure and pain the results are fundamentally the opposite. An organism must not do what it thinks will bring it pain, and should do what it thinks will bring it pleasure. Again, in each case pleasure and pain should stimulate the arousal system, but this time it is in anticipation rather than in actual sensation. The fact that anticipated pain is handled differently by the brain than current pain is has been shown by using brain imaging (Ploghaus, et al., 1999). The study used fMRI to show that pain and anticipated pain do activate some common regions (e.g. the medial frontal lobe), but that there were differences in activation as well, as pain and anticipated typically activated neighboring but distinct areas within a region. This, by the way, can be taken as evidence of the somatic marker hypothesis.

It should be noted that many emotional models do not explicitly include pleasure and pain, but do divide emotions into those that have positive or negative valence. The precaution of not including pleasure and pain directly may be due to the differences between cognitive and sensory pain (and pleasure). Pain researchers do not appear to be as reticent about the link. The definition of pain, according to the International Association for the Study of Pain is "an unpleasant

sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage" (Merskey, 1979, p. 250). The important point being that the emotional experience is intrinsic to the pain (Chapman, 1995; Craig, 1999).

It is at this stage that the basis for an emotional system can start to be seen. Having experienced pain in conjunction with some item, many creatures will begin to avoid the item as best they can. As Braitenberg points out with wonderful imagery (1986) this can be interpreted as a kind of hatred. Conversely if another item brings pleasure, creatures will try to be near the item. In more common terminology these behaviors are *avoid* and *approach*. From an adaptive perspective being able to predict pleasure and pain is even better than simply being able to detect pleasure and pain, but it must necessarily come after those simple abilities were established. Developmentally this is also true, as one must experience pain before one is able to predict it.

The step from simple pleasure/pain detection to the prediction of pleasure and pain is a large one and relies on several cognitive factors. First, there must be cognitive structure available that can be used to make predictions. Second, those predictions must be tagged somehow as being pleasurable or painful. Because pain and anticipated pain activate neighboring regions of the brain, Ploghaus et al. have speculated that this occurs through local interactions in those regions (1999). For example, someone putting their hand on a hot stove will experience pain. The combination of the cognitive structure active when the pain was experienced and the brain areas active because of the pain create a kind of associative link to other, neighboring, areas. Later when that same person thinks about putting their hand on the stove it is the neighboring areas that become active. This will recall the sensation of pain originally felt, but will not be equivalent to it. Presumably the difference is what allows the person to respond differentially in the two cases (e.g. reacting strongly in one case while doing nothing in the other).

The linkage between the cognition and the anticipation of pain is associative. Things that are experienced together in time are strongly associated cognitively. The strength of the association reflects factors such as the intensity of experience, as well as repetition. This association builds cognitive structure useful in prediction (Kaplan, et al., 1990). While the associative link to pain may not seem like a wonderful outcome, it serves a very useful purpose – namely providing a strong bridge between the cognitive system and the arousal system. In general, information- processing organisms do not need to "decide" to avoid painful things; they simply do because the combination of the pleasure/pain system, an arousal system, and predictive cognitive structure automatically makes it happen. In a sense the decision is left to evolution rather than to the individual and over the course of time this tradeoff has served individuals well. It is the strength of this response and its automaticity that makes emotions problematic to rationalists. If people were truly rational then all decisions would be made by weighing evidence and considering alternatives, but emotions dictate that many decisions are made on a different basis, one that favors fast action and safety. There are a number of flaws with the rational perspective. Among them is the fact that people rarely have the perfect information required for a true rational analysis. Probably even more important is the time required to make a "rational" decision. Emotional responses may be heuristic, but they are fast and the heuristics that they are based on have served innumerable generations.

Like the clarity and confusion mechanisms, the pleasure and pain equations are input driven. For sake of brevity, only pleasure is derived below. The inputs to pleasure include things that are inherently physically or cognitively pleasurable, as well as stimulation from the clarity system. These inputs may be seen in Tables 3 and 4 and can, to some extent, be considered additive. The point being that when one of the sources is particularly strong it will tend to overwhelm the rest and dominate. According to Kaplan (1991), in reality there are probably further levels of distinction based upon factors such as how similar the sources are, etc.

The pleasure inputs are derived according to the following algorithm. As previously stated, pain inputs would be similarly derived.

$$
\begin{aligned}
&\text{if}\left(\max(\rho_{+,1},\rho_{+,2},\ldots\rho_{+,n},C_{+}) > \rho_{threshold}\right)\\
&\quad \rho_{+} = \max(\rho_{+,1},\rho_{+,2},\ldots\rho_{+,n},C_{+})\\
&\text{else}\\
&\quad \rho_{+} = \min\left(\left(\textstyle\sum(\rho_{+,i}) + C_{+}\right),\rho_{threshold}\right)\\
&\\
&\text{where}\\
&\quad \rho_{+} = \text{pleasure inputs}
\end{aligned}
\tag{2}
$$

And, since Pleasure and Pain are, to some extent, mutually inhibitory, the algorithm for deriving the Pleasure and Pain inputs is given as:

$$
\begin{aligned}
&\text{if}(\lambda_{+}\rho_{+} > \lambda_{-}\rho_{-})\\
&\quad \text{if}(\lambda_{+}\rho_{+} > \rho_{threshold})\\
&\qquad P_{+} = \lambda_{+}\rho_{+}\\
&\qquad P_{-} = 0\\
&\quad \text{else}\\
&\qquad P_{+} = \lambda_{+}\rho_{+}\\
&\qquad P_{-} = \lambda_{-}\rho_{-}\\
&\text{else}\\
&\quad \text{if}(\lambda_{-}\rho_{-} > \rho_{threshold})\\
&\qquad P_{-} = \lambda_{-}\rho_{-}\\
&\qquad P_{+} = 0\\
&\quad \text{else}\\
&\qquad P_{+} = \lambda_{+}\rho_{+}\\
&\qquad P_{-} = \lambda_{-}\rho_{-}\\
&\\
&\text{where}\\
&\rho_{(+)} = \text{pleasure inputs}\\
&\rho_{(-)} = \text{pain inputs}\\
&\lambda_{(+)} = \text{susceptibility to pleasure constant}\\
&\lambda_{(-)} = \text{susceptibility to pain constant}
\end{aligned}
\tag{3}
$$

In our implementation of this model within the Agent domain, we assign pain and pleasure inputs according to Tables 3 and 4, respectively.

Table 3.

Factors Influencing Pleasure System

Pain Inputs	Value
I'm-hit **yes**	1
teammate-hit **yes**	0.5
danger **yes**	VAR
people-shooting **yes**	0.3
people-shooting-at-me **yes**	0.6
enemy-sees-me **yes**	0.5
moving-in-sight-of-enemy **yes** mission-in-jeopardy **yes**	0.25 0.30
communication-effective **no**	0.25
teammate-killed **yes**	0.5
high-enemy-to-friendly-ratio **yes**	0.3

Table 4.

Factors Influencing Pain System

Pleasure Inputs	Value
teammate-hit **no**	0.2
mission-accomplished **yes**	VAR
subgoal-accomplished **yes**	0.15
communication-effective **yes**	0.1
danger-passed **yes**	0.1
enemy-disabled **yes**	0.15

Arousal. Increased arousal has a number of well-studied effects on cognitive factors such as memory and attention (D'Ydewalle, et al., 1985; Hebb, 1972; Milner, 1991). Arousal is sometimes mistakenly linked to optimum functioning. The idea being that increased arousal leads to a kind of cognitive focus or sharpness. This is true to a point; however, as expressed graphically in Figure 3, Revelle and Loftus (1990) have shown that the link between arousal and performance forms an inverted-U curve such that the optimal level of arousal is neither too high nor too low. Thus, instead of thinking of arousal in terms of optimum performance, they suggest a better association for arousal is emotional intensity. As such, at median levels of arousal this intensity can indeed result in cognitive focus, but when arousal is too high emotions run towards the panic side.

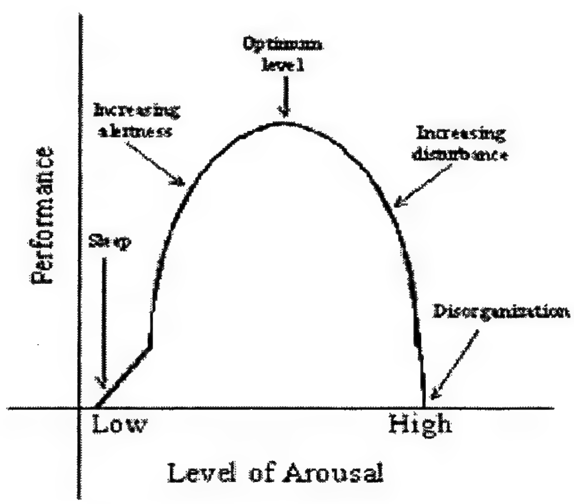

Figure 3. Generalized effects of arousal on performance.

To model increases in arousal, we use a model like that offered in Kaplan et al (1991):

$$\Delta A = S_A((A_B - ABS(A_B - A)) * ArousalInput) \quad (4)$$

where

$$ArousalInput = (\alpha_i + P_+ + \Delta P_+ + P_- + \Delta P_-)/3$$

A = arousal level

S_A = sensitivity to arousal constant

A_B = baseline arousal constant

α_i = arousal inputs

P_+ = pleasure value

P_- = pain value

In this case, pleasure and pain values would include quantities derived from the application of algorithms in (2) and (3) and an arousal input stimulation factors such as those shown in Table 5.

Table 5.

Factors Directly Affecting the Arousal System

Arousal Inputs	**Value**
loud-noise **yes**	0.3
dangerous-location **yes**	0.3
warning-message **yes**	0.25
move-to-cover-and-concealment **yes**	-0.2
enemy-moves-away **yes**	-0.2
OK-message **yes**	-0.1
Support-coming **yes**	-0.1

To represent the subsequent recovery of arousal to its equilibrium, we use

$$A_{Decay} = A_R(A - A_B) * (A_B - ABS(A_B - A)) \quad (5)$$

where

A_R = arousal recovery rate constant

Finally, combining the input and recovery equations yields the value of Arousal, as shown in equation 6.

$$A_{t+1} = A_t + \Delta A + A_{Decay} \quad (6a)$$

Which expands to

$$A_{t+1} = A_t + S_A((A_B - ABS(A_B - A)) * (\alpha_i + P_+ + \Delta P_+ + P_- + \Delta P_-)/3) - A_R(A - A_B) * (A_B - ABS(A_B - A)) \quad (6b)$$

All of the equations represented in rule-based form according to the Soar syntax may be seen in Appendix A.

Symbolic Modeling Representing Cognition

Although the emotional model as we have described it is basically connectionist, we have implemented it in the context of a rule-based system developed for the Soar architecture. This is because rule-based agents are simply more sophisticated and have a wider range of capabilities than their current neural network counterparts. The task we have chosen – that of a Special Forces team operating in enemy territory – is also appropriate because it is highly emotionally charged and fast decisions often have life or death consequences. The difficulty of implementing our emotional model in Soar, on the other hand, is that in some ways the model does not map cleanly to a rule-based system. The following sections describe the implementation of the emotions model in the rule-based system, SOF-Soar. The first section introduces the Soar architecture; the next section presents the SOF-Soar model and scenario we used for prototyping our concept; and the third section describes how the SOF-Soar model had to be parameterized in order to operate with the emotions model described in the previous section. This section is detailed according to how we implemented the model in general architectural terms, and then according to the implementation details specific to the SOF mission.

Soar

The original purpose of the Soar architecture was to support the development of intelligent systems that could use many different problem-solving methods. Soar quickly evolved to include integrated representations and methods for problem solving, planning, learning, and interaction with complex, dynamic environments. Many of the design requirements that fed into Soar came from early work on modeling human problem solving, by Allen Newell and Herb Simon (Newell and Simon, 1972). Soar has been used to develop many computational models of human problem solving and learning. All of these models share the same memory structure, type of task decomposition, task processing, and learning structure. Based on initial successes in modeling human behavior, Allen Newell proposed Soar as a candidate "Unified Theory of Cognition" (Newell, 1990). Different systems developed within Soar have successfully modeled a wide variety of human behavior relevant to this research (Rosenbloom, et al., 1993), and such work represents one of the few attempts to model a wide variety of psychological effects using a common software architecture.

A key component of all Soar models is that all activity is cast as a succession of decisions involving operators and goals. The decisions are based on an internal representation of the current situation, which is built up based on realistic simulated sensors. To make a decision, a Soar system performs a parallel retrieval from long-term associative memory (implemented as a very fast rule-based system) to get preferences for selecting the next "operator". An operator might represent an action as simple as "push the missile firing button", or as complex as "intercept a bogey". The retrieved preferences are analyzed, and a decision is made for the current best operator.

Once the current operator is selected, long-term memory is again consulted (via rapid, parallel, retrieval rules) to carry out the operator. If it is a simple operator, this will result in either a new output command being sent (which involves controlling simulated flight controls, weapon systems, sensor controls, or communication), or some changes to the internal state of the system (such as recording that a bogey is now being classified as a hostile bandit). If the selected operator is complex, such as intercepting a bandit, it will become a goal to be achieved through decomposition into one or more sub-operators. This activity can recurse, leading to the dynamic construction of an active goal hierarchy.

The long-term memory, together with the goal hierarchy, provides a smooth integration of reactive control and goal-driven behavior. Thus, the system quickly responds to changes in its environment, as it also selects new operators based on active goals.

Soar is written in C and has been ported to several operating system architectures (many UNIX variants, MacOS and Windows 95/NT) and C code development environments.

Special Operations Forces Behavioral Model

The Soar behavioral model used to evaluate the emotions model was Special Operations Forces (SOF) Soar. This task involves a 6-man team inserted deep within enemy territory. Once inserted, they travel anywhere from 20-50 km to the Objective Rally Point, they split into three 2-man teams (i.e., two 2-man observation teams, and one 2-man radio team).

Seeking cover and concealment, the observation teams set up near the designated Objective Observation Area and report back to the radio team via wireline radios when an appropriate objective has been sighted. The radio team will convey the essential elements of the observations back to the base using SATCOM or other secure channels. This is done at separate Transmit Sites, which are away from the observers and change after each transmission. At the conclusion of the mission, they will make their way to the Pickup Zone for exfiltration.

Scenario Parameters. Figure 4 shows an example of how the physical mission parameters may be arranged. In this mission, there are 5 types of critical points and areas: Drop point, Rally point, Objective and Transmit point, and Pickup point.

- **Drop point**: the location where the unit begins the mission (unless the Stationary option is selected).
- **Rally point**: the location where the unit moves to from the Drop Point. The Team leader and communications specialist remains at this location while the unit conducts its observation.
- **Objective**: the area to be observed.
- **Transmit point**: the location where the Radio team moves to in order to transmit an observation report to home base.
- **Pickup point**: the location to which the unit moves after completion of the recon and from which they are extracted. In this case, extraction simple means the entities disappear from the simulation.

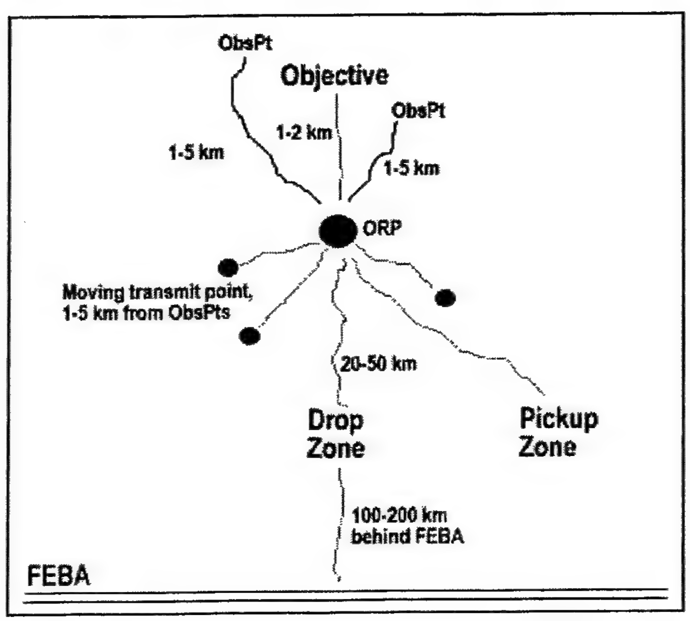

Figure 4. An example map of long range reconnaissance mission parameters.

Other important mission parameters include:

- **Start time**: Time at which the unit begins to execute the mission from the drop point.
- **Stop time**: Time that the unit departs their observation points.
- **Observation criterion**: Used to specify observation criteria of entities to be observed. Can be parameterized according to Type, Force, Size, Heading, and Heading Variation.
    - ***Type***: The type of target in which the unit has specific interest (e.g., AAA, APC, artillery, command, FWA, human, missile, RWA, ship, supply-truck, tank, vehicle).
    - ***Force***: The alignment of the force the operator wishes to observe (e.g., Red, Blue, Green and Any)
    - ***Size***: Restricts reports to organizations of the specified size or larger.
    - ***Radius***: The units will report contacts within this distance of their location (in meters).
    - ***Heading***: Reports only those units moving in this direction (degrees).
    - Heading Variation: Allows for the operator to specify a variation, in degrees, from the primary direction.

Some of these parameters may be changed in mid-mission. Typically, this would happen if the mission was jeopardized for some reason or perhaps not evolving as initially expected. In these events, it would be possible to activate the following changes:

- **Go Home**: causes the unit to go to the Pick-up point, by way of the Objective Rally Point (ORP).
- **Change commit criteria**: causes the unit to change its observation criteria.
- **Change pickup point**: Must be changed prior to the unit's arrival at the original pickup point.
- **Change the ORP**: Must be changed prior to the unit's arrival at the original ORP.
- **Change observation point**: Must be changed prior to the time the unit departs the ORP.

A run-time screen shot of our prototype may be seen in Figure 5, which shows the SOF Agents running in JSAF, an Agent's Soar Interface Panel, and an Agent's Emotions Interface Panel. The Soar Interface Panel enables operator control of an agent and communicates agent's decisions and actions. The Emotions Panel Interface is used to monitor the agent's emotional sub-systems and responses. To evaluate our system, we enhanced the SOF Agents with emotional responses, given some range of triggers. For example, we could alter the scenario by allowing detections, engagements, injuries, etc.

Scenario Enhancements. To make the existing SOF-Soar model more robust for experimentation with Emotions Model, a number of enhancements to the baseline Long Range Reconnaissance mission were required. Mostly, these enhancements were related to the addition of actions that the SOF-Agent could perform in the event of enemy contact. The rules developed to model possible reactions that a SOF-Agent might have in an emotionally charged scenario may be seen in Appendix B. A pseudo-code and natural language interpretation of these rules is provided in Appendix C, and the pseudo-code and natural language interpretation of the scenario assumptions is provided in Appendix D.

Implementation of Emotions in SOF-Soar

The two key parts of the symbolic component are an appraisal system and a response system. The response system accepts information from the connectionist model to influence behavior in a variety of ways. The implementation and natural language interpretation of the Emotions model, per equations 1 – 6b, may be seen in Appendices E, F, and G.

Cognitive Response to Emotion. The output of the emotional system includes numeric levels of arousal, pleasure, and pain. Arousal also creates a focusing of attention in the agent's higher level cognitive processes. The Soar architecture includes two types of memory (working memory and long-term memory), and arousal influences attention to elements in both of these memories.

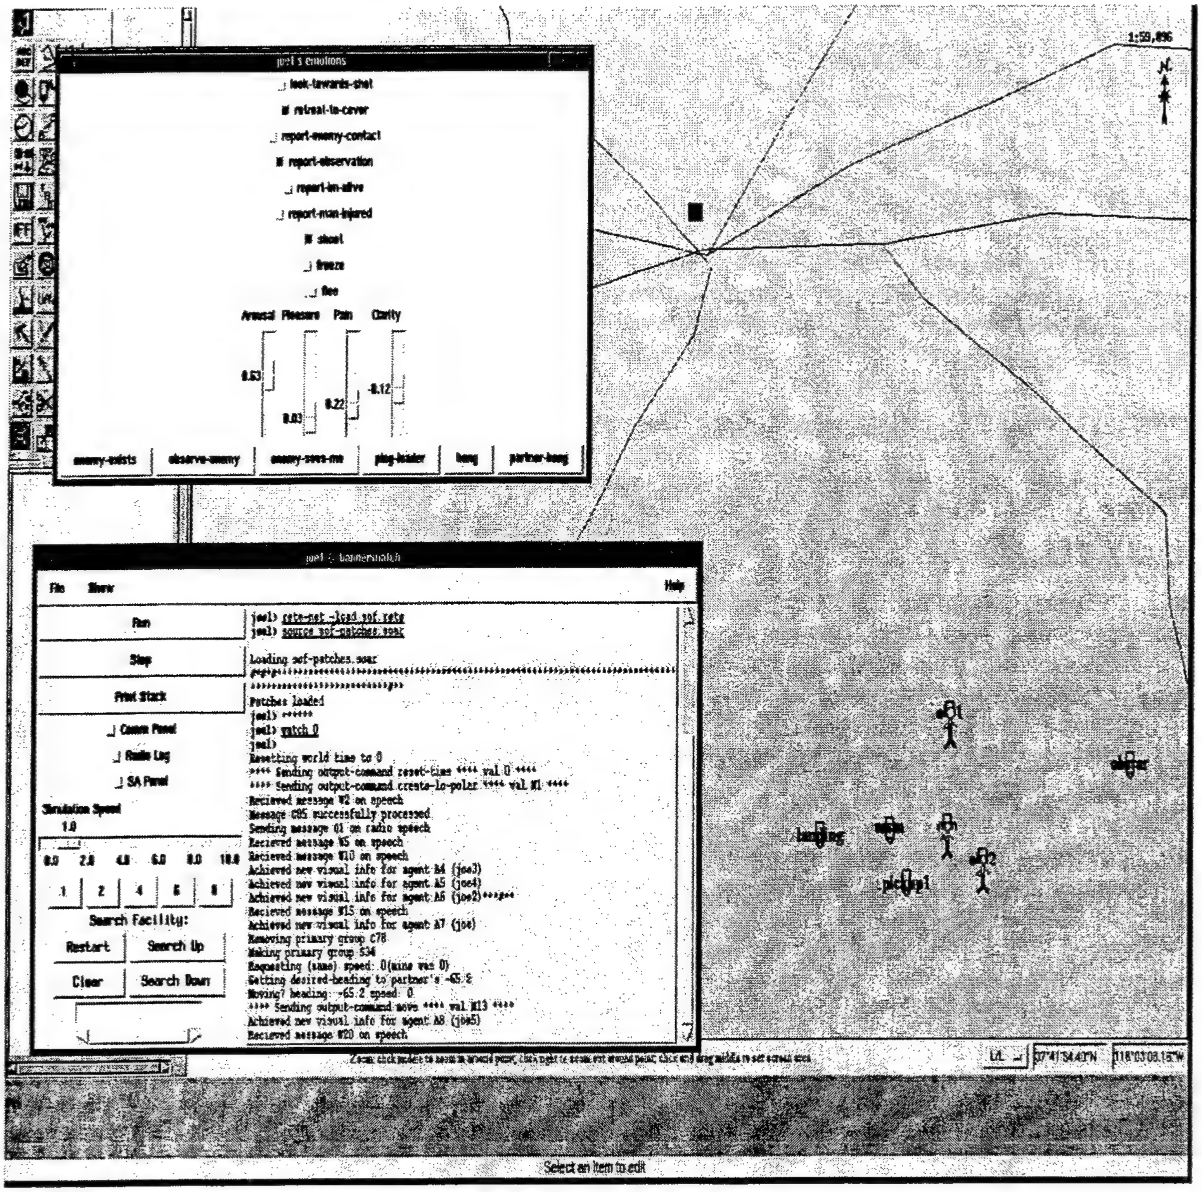

Figure 5. Runtime screen shot of emotional SOF Soar agent.

In working memory, arousal has similar consequences as it has for the agent's perceptual systems. That is, arousal tags particular working memory elements that are associated with the arousing stimulus. The "emotional" version of the SOF agent has extra preference knowledge (for guiding the temporal order of decision-making actions) that prefers attending to such tagged elements. Thus, the agent tends to focus its attention on the active concepts that are perceived to be relevant to the current level of arousal. In low-arousal situations, the agent's normal preference knowledge makes decisions with less regard for any particular focus of attention.

In long-term memory, arousal focuses attention on retrieving well-rehearsed chunks of knowledge with which the agent is very familiar and experienced. As discussed above, humans appear to prefer the comfort of using well-rehearsed knowledge in high-arousal situations. A consequence of this is that, when arousal is high, the intelligent agent will tend to revert to default patterns of behavior that are different from what the SOF agent would exhibit in more calm and deliberative situations. The default patterns of behavior vary across agents, depending upon personality (which we implement as different levels of susceptibility to emotional factors) as well as background knowledge. Our version of the SOF agent tags long-term knowledge with arousal thresholds, which filter the ability to retrieve particular chunks of knowledge. The Soar rules that model these filters are shown in Appendix E and the subsequent selection of actions are in

Appendixes F and G. These rules influence the SOF-Agent's decision-making process such that the Agent's reaction is selected according to its level of arousal and a global prioritization scheme.

The emotions subsystem also informs the cognitive component with current levels of pain and pleasure associated with arousing stimuli. Our implementation of the responses to pain and pleasure again take advantage of the existing preference mechanisms within Soar. Pleasure and pain most directly indicate to cognition whether particular types of decisions should tend to be pursued or avoided, depending on the valence associated with their outcomes. Soar's basic decision cycle implements a continuous series of decisions about "what to do next". The chosen actions may include creating new goals, recording new assessments of the environment, interpreting perceptual elements, predicting outcomes of actions or situations, or initiating external actions. In normal, "non-emotional", reasoning, a model developed within Soar typically includes preference knowledge to choose a single "next action" in the context of a complex system of situation interpretations and goals. For the most part, in an expert system, this preference knowledge encodes that rational selection of actions that maximize the chance of achieving goals. In the emotional SOF agent, any of the concepts or percepts involved in these decisions may be tagged with particular levels of pleasure or pain. Thus, by associating pleasure with goals, we can recast the notion of "always try to achieve the goal" to "always try to achieve pleasurable outcomes". In this manner, constructs that would normally be goals for a non-emotional system may become tagged with painful associations that cause a particular emotional agent to avoid pursuing them. Likewise, courses of action that become associated with high levels of pleasure may capture the attention of the emotional agent, particularly in high-arousal situations.

Appraisal. As the agent monitors its progress and the progress of its teammates, the appraisal system signals events that feed into the connectionist system, such as goal failure and achievement, and interpretations of the environment that cohere (suggesting clarity) or confound (suggesting confusion). The appraisal system emphasizes relevance to the system's current planning goals. For example, the agent should only be concerned about knowing the precise location of a particular enemy tank if that piece of knowledge is germane to the agent's current plan. Our efforts use appraisal rules to include low-level goals monitoring, as well as situation assessment (to gauge clarity and confusion), and appraisals in this consist of signals to the emotional system to adjust pleasure, pain, clarity, or confusion. The clarity and confusion appraisals are implemented in a more general fashion, while other appraisals take advantage of domain specific information. These appraisals ultimately result in adjustments that depend on the initial state of the connectionist network, as well as on a number of parameters dictating the strength of changes to the system.

Interfaces and Interactions between Emotion and Cognition

The implementation of the emotional network includes parameters that summarize a particular agent's base susceptibility to pain, pleasure, clarity, confusion, and arousal. Various settings of these parameters specify a "personality space" along the various dimensions of emotional effects. Thus, while cognitive appraisals may signal simple adjustments to each of these attributes, the actual effect of the adjustments will depend strongly on the personality profile of the agent. In turn, as mentioned above, the external effects of the attribute values will depend strongly on the agent's background knowledge, perception, and current understanding of its environment.

Impact of Emotions on Cognition

The primary way that the emotional system interacts with the cognitive system is through arousal. Arousal is sometimes even used interchangeably with emotional intensity. One reason for this is a study by Schachter (1962) where patients were injected with adrenaline and asked to describe their feelings as euphoria or anger. When subjects were induced to discount the injections as the cause of their aroused state they described this state in emotional terms appropriate to the experimental cues they received. Whereas pleasure/pain and clarity/confusion

work to detect events of importance to an agent, the arousal system functions as a kind of interface between the emotional and higher cognitive systems. Increased arousal has a number of well-studied effects on cognitive factors such as memory and attention (Hebb, 1972; D'Ydewalle, et al. 1985; Milner, 1991).

Memory and attention are the cognitive components most affected by changes in arousal. Highly aroused people are likely to fall back on well-learned knowledge and habits, even when they might have more relevant knowledge available. Evolutionary theorists have explained this by relating familiar events with safety. Because highly arousing situations are of extra importance, it is considered best to rely upon what has worked in the past, rather than on things that are not as well tested. In addition, arousal has predictable effects on learning (D'Ydewalle, et al. 1985; Milner, 1991). As befits important situations, people are more likely to have strong long-term memory for experiences that occurred while they were highly aroused than for situations where they were not.

Impact of Cognition on Emotions

Differential levels of arousal constitute the primary output of the emotional system (though the valence of the emotions is equally important). Inputs to the emotional system, on the other hand, can come in a wide range of forms from the direct sensation of pleasure and pain to highly cognitive situational assessments.

Many emotional inputs are essentially perceptual. Pleasure and pain can be sensed directly, and some perceptions can be either inherently arousing (or calming). It is even the case that perceptions can be inherently clear or confusing. Kaplan and Kaplan (1982), in analyzing human preference for visual scenes, found that one of the critical dimensions was "legibility," a concept that maps directly to the clarity/confusion distinction. Highly legible scenes typically have a great deal of recognizable structure whereas illegible scenes are seen as chaotic. The Kaplans have interpreted the preference for legibility as stemming from the fact that people feel more comfortable with such environments because they expect to be able to perform more effectively within them. All of the direct inputs to the emotional system are the result of evolution and reflect hard-wired rules for important events. Experience provides additional information of what is important, and therefore the cognitive system feeds back into the emotional system. One way that this occurs is through memory and through using cognitive structure to make predictions. For example, when experiencing pain, memory codes the experience as painful. Then, anticipating a similar event will elicit cognitive pain, and therefore increase arousal.

While predictions provide concrete assessments of what is to come, sometimes there is not enough information available to make a detailed analysis. In such cases the likelihood of making good decisions can be based on the perceived level of competence within the environment. The simplest way to do this is to continually compare what one expects to happen versus what actually happens. Predictions that come true reflect an excellent world model combined with good information. In such situations performance is likely to be good, and this clarity should engender positive feelings. By contrast, when predictions either cannot be made for a lack of information, or do not come true, this is a highly dangerous situation that is likely to be unsettling and upsetting because it means that good decisions are unlikely.

Figure 6 shows the overall integration and interactions between these sub-components. Essentially, the SOF-Soar Agent interacts with the environment through the Soar-ModSAF Interface (SMI) where it accesses information on the environment. This could include situational awareness data from reports or perceptions. This data is stored in the agent's working memory and accessed by the emotions model through the emotions interface. Thus, the emotions interface only interacts with the SMI and the SOF-Soar Agent; it does not directly interact with the simulation environment. Based on situational awareness and current information in working memory, the emotions model computes new values for confusion/clarity, pain/pleasure, and arousal. Then, these values are returned to the agent's working memory where they are coupled

with the context and used to propose, filter, and select the agent's next action. This cycle (i.e., the cycle between working memory, emotions interface, behavior moderators, and long-term memory) will continue until all activity in working memory has acquiesced. Once this happens, the SOF Soar Agent will select a motor command (some action) and send this information to the simulation environment for processing.

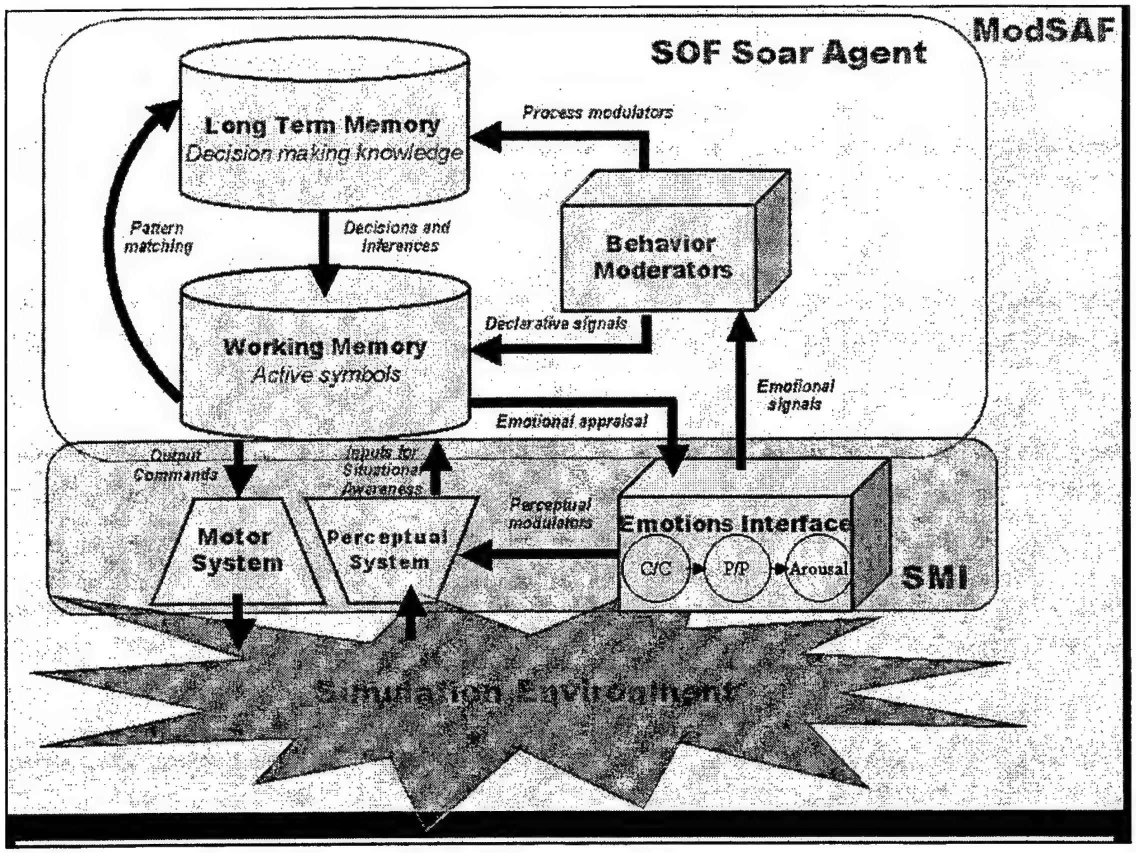

Figure 6. Architecture of emotional SOF-Soar system.

Emerging Behavior

The interaction amongst the emotional subsystems and the interactions between the emotional system and cognition form the foundation of a model that exhibits emerging behavior. Emerging behavior, simply defined, is an attempt to understand high-level behavior from low-level rules; for example, how the simple rules of Darwinian evolution lead to high-level structure. Emergent behavior occurs when a system produces unexpected behavior according to non-linear interactions amongst the system's sub-components. That is, emergence refers to the appearance of higher-level properties and behaviors of a system that are not directly deducible from the lower-level properties of that system (Ilachinski, 1996).

Individual-based models (IBMs) are models that show evidence of emerging behavior in that they are simulations based on the global consequences of local interactions of members of a population. These models can also referred to as entity-based or agent-based models or simulations and they typically consist of an environment or framework in which the interactions occur and some number of individuals (e.g., plants and animals in ecosystems, vehicles in traffic, people in crowds, or autonomous characters in animation and games) defined in terms of their

behaviors (procedural rules) and characteristic parameters that are tracked through time. This stands in contrast to modeling techniques where the characteristics of the population are averaged together and the model attempts to simulate changes in these averaged characteristics for the whole population.

The model proposed in this investigation makes use of the concept of emerging behavior in two ways. First, the emotion model, in isolation is constructed around this concept. That is, the emotion experienced by an agent is ultimately the result of a combination of factors (e.g., arousal, pain/pleasure, and clarity/confusion) that interact over time depending on triggers in the scenario. To the extent that these triggers can involve other agents that can be assigned a unique temperament and emotional state, this becomes another type of emergence. That is, in this latter case, the order of emergence now depends on an additional, higher-ordered set of behaviors emerging from actual entity interaction (as opposed to isolated triggers causing behavior through the interaction of the emotional model substructures). Thus, we have a within-agent emerging behavior pattern (i.e., the interaction of the emotional model substructures depending on the environment) and a between-agent emerging behavior pattern (i.e., the interaction between agents, both/all of whom can have unique emotional states and temperaments). Taken together, the whole creates a fairly complex feedback system, in which the resulting external behavior would be very difficult to predict analytically. This justifies the approach of building these models within executable intelligent agents, so that the resulting behaviors and emotional effects can be characterized empirically. The fundamental motivation of this research is to explore high-level emergent behaviors arising from various low-level "interaction rules."

Experimentation

Once an emotions model has been established, we need some organizing framework within which to exercise it. After all, different people have different reactions to the same situations (i.e., emotions and emotional responses are unique to individuals). Such differences can be thought of as an "emotional style" or temperament. We use the body of research in emotions and temperament to develop the bounds of an experimental design region for testing our model. Adopting such a framework allows us to distinguish between the characteristic way a person experiences emotions (emotional style) and the way those emotions are realized (emotional content). In this framework, for example, referring to someone as frightened or angry would involve referring to the content of a person's emotional experience. On the other hand, referring to someone as hot headed or stoic would involve referring to the style with which a person may be inclined to experience emotions.

The first section under experimentation explains the concept of using personality space to bound our experiments and reviews the literature germane to the development of this framework. The next section presents validation methodology static test cases used to evaluate this framework through cases of increasing complexity. The follow-on section, evaluates the more complex, dynamic and interactive cases. Thus, the general form to our testing procedure is to implement the simplest cases first and then allow the more complex cases to emerge as a result of behavior that has been previously established and evaluated.

Personality as Experimental Design Space

Because it is our intent to examine how emotions influence decisions differently for different type of people and in different circumstances, a critical part of our experimental plan is to test a variety of identical scenarios using SOF Agents with different emotional templates. This being the case, it is important to relate our model to traditional notions of emotional state and to supplement this with an empirical investigation that considers a variety of "personality types" that vary along emotional lines. Based on the connectionist structure of the emotional model, we use the connections between components as parameters that partially determine personality type. For

example, one parameter in the arousal component can determine an individual's sensitivity to becoming aroused. The psychological literature has long theorized that the critical factor that distinguishes introverts and extroverts is the relative susceptibility to becoming aroused (e.g., Eysenck & Eysenck, 1985, D'Ydewalle, et al. 1985). Thus, the ability to adjust this parameter in the arousal component of the emotion model provides a mechanism to represent the introvert/extrovert dimension of personality. Individual differences in susceptibility to arousal and other emotional attributes lead to distinctive decision making profiles that can lead to crucial performance differences in emotionally charged situations. We should stress that we are not developing a full-fledged theory of personality. Rather, it is our contention that a great deal of a person's personality can be explained by their emotional profile as exemplified by parameters such as susceptibility to being aroused. In particular the three major components of our emotion model appear to map onto three of the five personality traits identified as part of the "Big 5" theory (Digman, 1990; Costa and McCrae, 1995), currently the leading personality descriptive system. We link the three emotional mechanisms in our model to the dimensions of introversion/extraversion, positive/negative emotionality, and preserver/explorer in the Big 5 model. We postulate that "negative emotionality", the amount of a stimulus required to elicit a negative response, is defined by a person's sensitivity to pleasure and pain. The final trait that we have identified is "openness." The Big 5 theory divides people along this category into "preservers" and "explorers," with the key difference being how open they are to new and different experiences. It is our assumption that preservers have a low tolerance for confusion and that perhaps explorers can find clarity even in relatively novel situations. One advantage of this framework is that standard personality tests exist to divide people along each of these categories. Of course, our model will not have any sort of explicit output that marks an agent as a "preserver" or an "introvert". The purpose of the experiments will be to correlate agent behavior in different situations to the expected behavior of people with these attributes. The remaining two dimensions of the Big 5 theory are not related to emotions (at least at the level of our emotional model), but they are also not inconsistent with our architecture.

Returning to Figure 2, we can see that the various parts of our emotions model are linked together. Clarity brings pleasure, which in turn brings arousal, etc. Each of these major pieces represents a source of individual difference and these differences can be explained by using the links between the pieces. If the figure were to be viewed as a neural network (or a vast simplification of a network) the links would be connections between nodes. Individual differences in this case would simply occur as differently weighted links in the model. For the purposes of this report it may be simpler to think of each piece individually and consider the individual differences to arise out of a difference in susceptibility to each. For example, some people may be highly susceptible to experiencing pain. In the model this could be expressed by having strong links into the pain portion of the model. To simplify our discussion we will consider three basic type differences and go on to show how those differences might be expressed as personality types.

The first difference has to do with the susceptibility to becoming aroused. Individual differences in this case mean that one person might remain calm in the face of horrific circumstances and even personal pain while another might become highly agitated at the slightest disturbance.

The second difference has to do with pleasure and pain. We do not rule out the possibility that a single individual might be differentially sensitive to pleasure versus pain, but given the lack of direct evidence for it and the nature of this discussion, we will assume that different people experience pleasure and pain at different thresholds. There is even some evidence that links these different thresholds to eye color (Rosenberg & Kagan, 1987). A person more sensitive to pain is more likely to focus on negative events, while a person who is less sensitive might be more stoic. It is worth noting at this point, that pleasure and pain are mutually inhibitory in our model.

The final difference has to do with clarity and confusion. The individual differences marking clarity and confusion are less obvious than in the other cases. This is because they are confounded with the kinds of cognitive structure that an individual generates. Fast learners may

not become confused as often as other people due more to how they learn, than to how prone they are to being confused. Nevertheless, it is reasonable to assume that different people have different thresholds in become confused or in experience clarity. On the one hand, an easily confused person might shy away from the novel or the complex because of the negative feelings associated with being confused, on the other some people may be able to find clarity where others see only chaos.

At this point is worth discussing the clarity/confusion continuum in a little more depth. Unlike pleasure and pain, clarity and confusion are not polar opposites. Confusion probably is at one end of the continuum, but its opposite is familiarity or boredom. As cognitive structure develops it goes from the confusing and novelty into clarity and discovery. At the level of clarity, a cognitive structure still has some potential for exploration, but with the safety of competence. At this point further experience has little to add to shaping the structure, but can only serve to make it leaner and more efficient in a kind of cognitive chunking. For example, as people learn to drive they go from needing to think everything through in their heads, to responding automatically and without thought. Processing with such structures is more efficient and demands less capacity. Evolution, however, has provided a mechanism to push the boundaries of what we know. Beyond clarity is boredom. As cognitive structures become hyper-efficient chunks they no longer generate intense enough activity to drive clarity because they only use a small part of attention. Boredom pushes people to experience the novel once more and to begin to build more cognitive structure. From an evolutionary perspective, the well learned is safe. It is better to be facile in more environments, but it is safest to learn them at a pace that allows for the retreat to the known. Hebb framed these issues in terms of play, and the need for excitement (1972).

The dominant current theory in personality research is generally referred to as the "Big 5" theory (Digman, 1990; Costa and McCrae, 1995). Statistical analysis has generally shown that there are reliably five orthogonal personality dimensions. In practice there is general agreement about three of the dimensions, and a little less about the final two dimensions. Another prominent personality theory, Psychoticisms/Extraversion/Neuroticism (PEN), proposed by Eysenck (1991) proposes three dimensions. We will focus on two of the three dimensions that personality theorists seem to agree on, and attempt to show that each dimension can be understood as an individual difference in the emotional model we are presenting. The three factors we will focus upon are generally called extraversion, negative emotionality, and openness. The two we will spend less time with are agreeableness and conscientiousness.

Extraversion

Extraverts are typically categorized as people who prefer to be with other people. Extraverts are outgoing and assertive and sometimes are described as craving excitement. Introverts on the other hand prefer to be by themselves and are often described as reserved. Introverts tend not to seek excitement. Within this category there is some research that indicates that extraverts also tend to have more positive affect than introverts (Gross, et al., 1998; Carver, et al., 2000).

The arousal literature has long made the case that extraverts are more difficult to arouse than introverts. The excitement needed by extraverts is to generate arousal. Going back to the notion that graphing performance against arousal generates an inverted U, extraverts tend to be on the lower left end of the U and need excitement to push themselves to the middle. The stimulation provided by people and novelty can provide such excitement. Introverts, on the other hand, will tend to be pushed to the lower right of the U when in stimulating environments. Therefore they will tend to try and control for this by seeking quieter places and by being alone.

It is tempting to postulate that the source of excitement for extraverts is pleasure. This would tend to be supported by the studies linking extraversion and positive affect (Watson & Tellegen, 1985). However, this seems unlikely to be the case and another explanation is more plausible. Over the course of their lifetimes, extraverts will learn what environments they do well

in, and what environments they do not do well. In novel environments, especially those that are stimulus intensive, extraverts will expect to perform well and will develop a confident attitude accordingly. Introverts will not tend to have such positive experiences. Because they become overstimulated, they will not perform as well and will generally come to associate those feelings with such environments. The argument here is simply that extraverts will tend to have more positive associations with many kinds of environments because of their experiences. On their own and free of stimulation, they should not be any more prone to feeling good than anyone else. This conclusion is supported by studies that show exactly that extraversion correlates only very weakly to the magnitude of changes in positive affect for positive stimuli (Gross, et al., 1998).

Negative Emotionality

This dimension is differentiated by the strength of a stimulus required to elicit a negative response. At one end of the spectrum are people who are "resilient." They are described as calm, slow to discourage, and handle stress well. At the other end are people who are "reactive." These people are uneasy, quick to anger and embarrass, and do not handle stress well.

Since this factor is framed in terms of negative emotions, it is tempting to simply link it to pain. This is probably a mistake since pleasure and pain are so closely linked. Rather it is more likely the case that we simply pay more attention to pain, partly because pain is more important from an evolutionary perspective (since pain is equated to danger). Regardless, the central point is that people who experience pain more easily are more likely to focus on it in their daily lives. This factor is sometimes labeled "neuroticism" and is often related to anxiety. Our contention is that this anxiety comes as the result of low levels of pain experienced to a nearly constant degree. In many people this pain is easily ignored, but for people with a lower threshold the pain will be very real and they will come to associate it with most aspects of their lives. Viewed from this perspective neurotics may be responding to life in a very reasonable way. The same studies that showed only a weak link between extraversion and changes in positive affect, showed very strong correlations between neuroticism and changes in negative affect (Gross, et al., 1998).

Openness

This dimension is more controversial than the previous two and has been left out in three dimensional models such as Eysenck's (1981; 1991). This trait divides people into explorers and preservers. The traits themselves suggest the linkage to boredom and play. Preservers are described as interested in the here and now, preferring the familiar and conservative about change. Explorers, by contrast, daydream, are open-minded and prefer variety.

It should not be surprising that this trait is more controversial than others as it is probably muddied somewhat by the interplay of clarity and confusion with how people learn. Nevertheless, the clarity/confusion mechanism provides insight into the different types. Presumably preservers do not like new things because they abhor becoming confused. It is simpler to fit something into existing structure than to deal with learning new structure. Explorers, on the other hand, may be addicted to clarity and easily prone to boredom. Too much of the same thing will cause them pain, so they will seek the novel.

Agreeable and Conscientiousness

The final two dimensions of the Big 5 model are agreeableness and conscientiousness. Agreeableness is a category generally included in personality models and reflects the sources that people use to define correct behavior. This would appear to be a social factor and is thus, not mapped to our emotions model. Conscientiousness is defined in terms of goals. Again, this does not directly map to emotions in our model.

As demonstrated in Figure 7, our plan is to map out the space of "personality types" dictated by these dimensions. We will then instantiate particular agents that characterize typical and extreme portions of this space. Finally, we will run each agent in a suite of test scenarios designed to highlight situations in which particular types of emotional responses are beneficial and detrimental.

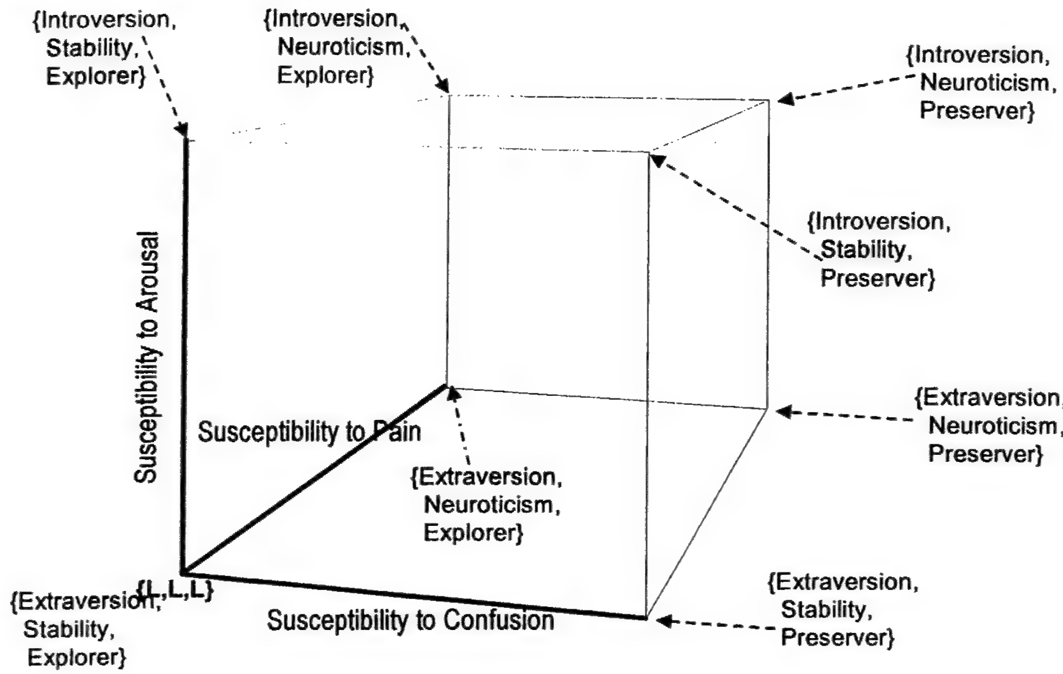

Figure 7. Personality space defines experimental region for emotions model.

Validation Methodology

Because behaviors of IBMs exist only in software and are based on nonlinear and multi-threaded process controls, they are difficult to verify through conventional, analytical methods (Riekel, 1995; Ropella, 2002). For example, the emotional SOF Agents discussed in this paper exist only in software. And, like other models of emerging behavior, because this representation is theoretically based and lacks data for conventional forms of validation, we adopt principles of the IBM community (Riekel, 1995; Ropella, 2002) in validating the model. Essentially, this approach considers validation an issue of deciding whether the model output meets the required performance standards according to the model's purpose. In the IBM community, this approach is known as "Operational Validation". Of course, another important part of this process is verifying the code that is, checking that model does what it is intended to do.

Riekel offers a number of strategies for verifying and validating the software accuracy of IBMs or other similar types of models based on principles of emerging behaviors. Central to these strategies is a clear experimental procedure, where testing is treated as a scientific enquiry with testers designing experiments, predicting the outcomes, and then running the code to compare the actual outcomes to the predicted outcomes. Verification and validation strategies offered by Riekel include: code reviews, spot checks of key model subcomponents, pattern tests, and systematic tests against an independent implementation.

To test the model presented in this research, we adopt a dual approach, measuring the within-agent emerging behavior pattern (i.e., the interaction of the emotional model substructures and cognition for an individual agent, given a personality type) and the between-agent emerging

behavior pattern (i.e., the interaction between agents, both of whom can have unique emotional states and temperaments). These tests are performed over a number of increasingly complex test cases. Initially, as a means of verifying model code, we consider static test cases to measure the within-agent patterns. These measures include verifying the emotional system's numeric output according to personality types, and verifying that the correct response is selected, given some state vector. All of these measures are compared against a manual simulation (implementation of equations 1 – 6a/b for each static test case) of each test case to determine accuracy. Follow-on tests evaluate the more complex cases measuring the dynamic behavior emerging from between-agent interaction. As in the simple case, these tests also systematically increase in complexity. The test cases are presented in the next section.

Testing the Static Case with Personality Framework

The general scenario for testing the static case with a personality framework is built around the two-man SOF team at the Observation Point. In all, there are 5 derivatives of this basic scenario, where each of these derivatives becomes more complex. The basic assumptions and state of scenario is specified in each scenario, along with the subsystem inputs that apply to the situation. Since these cases were all evaluated statically, an initial current arousal level of .501 was assumed. Thus, a comparison between "Next Arousal" value and .501, as well as between "Next Arousal" values across personality parameters is used to confirm that the "Next Arousal" value is consistent with expectations.

Scenario 1 – At Observation Point. In this scenario the two-man SOF Observation team is stationed at the Observation Point and there is no enemy in sight. The inputs to the emotional subsystem for this case may be seen in Table 6. For this simple static case, the results according to personality type are shown in Table 7.

Table 6

Subsystem Inputs Considered for Scenario 1

Emotional Subsystem	Inputs	Value
Arousal	dangerous-location yes	.3

Table 7

Next Arousal Value Resulting from Scenario 1

			NEXT AROUSAL
Extraversion	Neuroticism	Explorer	0.51347999
Extraversion	Neuroticism	Preserver	0.51347999
Extraversion	Stability	Explorer	.0.51347999
Extraversion	Stability	Preserver	0.51347999
Introversion	Neuroticism	Explorer	0.53842999
Introversion	Neuroticism	Preserver	0.53842999
Introversion	Stability	Explorer	0.53842999
Introversion	Stability	Preserver	0.53842999

Scenario 2 – At Observation Point and Enemy Sighted. In this scenario the two-man SOF Observation team is stationed at the Observation Point and a high number of enemies have been sighted. The inputs to the emotional subsystem for this case may be seen in Table 8. For this case, the results according to personality type are shown in Table 9.

Table 8

Subsystem Inputs Considered for Scenario 2

Emotional Subsystem	Inputs	Value
Confusion	unknown-value useful	.25
Pain	high-enemy-to-friendly-ratio yes	.3
Arousal	dangerous-location yes	.3

Table 9

Next Arousal Value Resulting from Scenario 2

			NEXT AROUSAL
Extraversion	Neuroticism	Explorer	0.53219249
Extraversion	Neuroticism	Preserver	0.53219249
Extraversion	Stability	Explorer	0.51971749
Extraversion	Stability	Preserver	0.51971749
Introversion	Neuroticism	Explorer	0.59456749
Introversion	Neuroticism	Preserver	0.59456749
Introversion	Stability	Explorer	0.55714249
Introversion	Stability	Preserver	0.55714249

Scenario 3 – At Observation Point, Enemy Undetected, and Shooting. In this scenario the two-man SOF Observation team is stationed at the Observation Point and detects shooting, although they have not seen enemy. The inputs to the emotional subsystem for this case may be seen in Table 10. For this case, the results according to personality type are shown in Table 11.

Table 10

Subsystem Inputs Considered for Scenario 3

Emotional Subsystem	Inputs	Value
Confusion	enemy-exists yes	.4
Confusion	enemy-sees-me yes	.2
Pain	people-shooting yes	.3
Pain	people-shooting-at-me yes	.6
Pain	enemy-sees-me-yes	.5
Pain	mission-in-jeopardy	.3
Pain	high-enemy-to-friendly-ratio yes	.3
Arousal	loud-noise	.3
Arousal	dangerous-location yes	.3

Table 11

Next Arousal Value Resulting from Scenario 3

			NEXT AROUSAL
Extraversion	Neuroticism	Explorer	0.56337999
Extraversion	Neuroticism	Preserver	0.56337999
Extraversion	Stability	Explorer	0.53842999
Extraversion	Stability	Preserver	0.53842999
Introversion	Neuroticism	Explorer	0.68812999
Introversion	Neuroticism	Preserver	0.68812999
Introversion	Stability	Explorer	0.61327999
Introversion	Stability	Preserver	0.61327999

Scenario 4 – At Observation Point, Enemy Undetected, Shooting, and Teammate Hit. In this scenario the two-man SOF Observation team is stationed at the Observation Point and teammate has been shot, although no enemy has been detected. The inputs to the emotional subsystem for this case may be seen in Table 12.

Table 12

Subsystem Inputs Considered for Scenario 4

Emotional Subsystem	Inputs	Value
Confusion	enemy-exists yes	.4
Confusion	enemy-sees-me yes	.2
Pain	teammate-hit yes	.5
Pain	people-shooting yes	.3
Pain	people-shooting-at-me yes	.6
Pain	enemy-sees-me-yes	.5
Pain	mission-in-jeopardy	.3
Pain	high-enemy-to-friendly-ratio yes	.3
Arousal	loud-noise	.3
Arousal	dangerous-location yes	.3

Table 13

Next Arousal Value Resulting from Scenario 4

			NEXT AROUSAL
Extraversion	Neuroticism	Explorer	0.56337999
Extraversion	Neuroticism	Preserver	0.56337999
Extraversion	Stability	Explorer	0.53842999
Extraversion	Stability	Preserver	0.53842999
Introversion	Neuroticism	Explorer	0.68812999
Introversion	Neuroticism	Preserver	0.6881299
Introversion	Stability	Explorer	0.61327999
Introversion	Stability	Preserver	0.61327999

Scenario 5 – At Observation Point, Enemy Undetected, Shooting, and I'm Hit. In this scenario the two-man SOF Observation team is stationed at the Observation Point and the agent has been shot, although no enemy has been detected. The inputs to the emotional subsystem for this case may be seen in Table 14. For this case, the results according to personality type are shown in Table 15.

Table 14

Subsystem Inputs Considered for Scenario 5

Emotional Subsystem	Inputs	Value
Confusion	enemy-exists yes	.4
Confusion	enemy-sees-me yes	.2
Pain	I'm-hit yes	1.0
Pain	people-shooting yes	.3
Pain	people-shooting-at-me yes	.6
Pain	enemy-sees-me-yes	.5
Pain	mission-in-jeopardy	.3
Pain	high-enemy-to-friendly-ratio yes	.3
Arousal	loud-noise	.3
Arousal	dangerous-location yes	.3

Table 15

Next Arousal Value Resulting from Scenario 5

			NEXT AROUSAL
Extraversion	Neuroticism	Explorer	0.58832999
Extraversion	Neuroticism	Preserver	0.58832999
Extraversion	Stability	Explorer	0.54674665
Extraversion	Stability	Preserver	0.54674665
Introversion	Neuroticism	Explorer	0.76297999
Introversion	Neuroticism	Preserver	0.76297999
Introversion	Stability	Explorer	0.63822999
Introversion	Stability	Preserver	0.63822999

Testing Dynamic Case

The general form to our testing procedure is to implement the simplest cases first and then allow the more complex cases to emerge as a result of the behaviors that have already been established and evaluated. For example, for the within-agent tests, each of the scenarios 1-5 (see Table 16) is used to statically evaluate the emotional system, in isolation. This means there

Table 16

Progression of Test Scenarios

Scenario Title/Variation	Description
1. At Observation Point.	Two-man Observation team is stationed at the Observation Point and there is no enemy in sight.
2. At Observation Point and Enemy Detected.	Two-man Observation team is stationed at the Observation Point and a high number of enemies have been sighted.
3. At Observation Point, Enemy Detected, and Shooting.	Two-man Observation team is stationed at the Observation Point, spotted enemy, and detected shooting.
4. At Observation Point, Enemy Detected, Shooting, and Teammate Hit.	Two-man Observation team is stationed at the Observation Point, detected enemy, and agent's teammate has been shot.
5. At Observation Point, Enemy Detected, Shooting, and I'm Hit.	Two-man Observation team is stationed at the Observation Point, detected enemy, and the agent has been shot.

were five independent scenarios for the within-agent tests. Thus, even in this simple block of static scenarios, a total of 160 tests are performed (emotional levels of each personality type for each scenario), assuming some baseline arousal level. Still focused on within-agent behavior, the next more complicated round of tests evaluates the first-order case where emotion and cognition interact. Again, this is accomplished by verifying model results with independently implemented manual simulations.

Once these initial results are verified through static tests, dynamic cases starting with scenario 1 progressing to variations of scenarios 4/5 are executed to record patterns in between-agent behavior. Differences in behaviors over these scenarios will be due to differences in agent's arousal level and how that impacts cognition, where these components of the model were previously verified in within-agent tests. In other words, for the within-agent static cases, an initial cognitive state was assigned, and then arousal was computed according to this state. Thus, the "complete" cycle between cognition-arousal was fulfilled, but never allowed to iterate more than once. In the dynamic case, however, where the scenarios in Table 16 are actually "progressions" of a single scenario, the cognition-arousal computational cycle is feeding back on itself. This allows us to watch an individual agent's behavior over time. Additional differences, that may still be isolated through this latter set of tests, are the result of how the behaviors of one agent can impact the emotional intensity, and hence response, of another agent. This test, however, has not been performed.

As stated in the previous section, the primary focus of the validation effort is to develop a sense of the model's utility, given its purpose. The objective of this research, as communicated in abstract, was to make the decision-making process of complex agents less predictable and more realistic, by incorporating an emotions model. The following sub-sections present results pertaining to both of these cases. In the first sub-section, we demonstrate the model's utility by presenting scripted output of simple test scenario. Next, in the second sub-section, we report on means of determining how this system reduces the predictability of an agent's behavior.

Output behavior of one example of a simple dynamic case is scripted in Table 17. This test case contrasts the behavior of Agent assigned emotional style of <Extravert, Stability, Explorer> with an Agent assigned emotional style of <Introvert, Neurotic, Preserver>. In this scenario, the SOF Agents are at the Observation Point and detect enemy, the objective on which Agents should report. Up to this point, even though the Agents exhibit differences in the values of

their emotional parameters, the Agents propose and select the same reactions based on the same world events. The next event, "Enemy-Sees-Me" causes Agent2 to propose one more action ("flee"). However, both Agents choose to "Retreat-to-Cover". During the next event, "Partner-Shot", behavior of the two Agents starts to diverge. That is, Agent1 chooses to report the injury, whereas Agent2 continues to seek cover. Lastly, as the "Shooting" (final event) continues, Agent1 remains active in seeking cover, whereas Agent2 "freezes", in essence rendering him useless in the rest of the scenario.

Table 17

Example of Simple Dynamic Test

Event	Agent1(ESE)	Agent2(INP)
At ObsPt	Arousal = .54 Pleasure = .2 Pain = .0 Clar/Conf = .15 **Proposed action(s):** no change **Selected action:** no change	Arousal = .77 Pleasure = .59 Pain = .0 Clar/Conf = .35 **Proposed action(s):** no change **Selected action**: no change
Observe Enemy	Arousal = .54 Pleasure = .15 Pain = .06 Clar/Conf = .09 **Proposed action(s):** report-observation **Selected action:** report-observation	Arousal = .77 Pleasure = .43 Pain = .14 Clar/Conf = .21 **Proposed action(s):** report-observation **Selected action:** report-observation
Enemy-Sees-Me	Arousal = .54 Pleasure = .03 Pain = .19 Clar/Conf = -.03 **Proposed action(s):** retreat-to-cover, report-observation, shoot **Selected action:** retreat-to-cover	Arousal = .76 Pleasure = .07 Pain = .47 Clar/Conf = -.07 **Proposed action(s):** retreat-to-cover, report-observation, shoot, *flee* **Selected action:** retreat-to-cover
Partner-Shot	Arousal = .62 Pleasure = .00 Pain = .22 Clar/Conf = -.12 **Proposed action(s):** retreat-to-cover, shoot, *report-man-injured* **Selected action:** *report-man-injured*	Arousal = .93 Pleasure = .00 Pain = .62 Clar/Conf = -.28 **Proposed action(s):** retreat-to-cover,shoot **Selected action**: *retreat-to-cover*
Shooting	Arousal = .87 Pleasure = .0 Pain = .27 Clar/Conf = -.39 **Proposed action(s):** flee, *retreat-to-cover* **Selected action:** *retreat-to-cover*	Arousal = .98 Pleasure = .0 Pain = .63 Clar/Conf = -1.0 **Proposed action(s):** flee, *freeze* **Selected action:** *freeze*

Discussion

This section summarizes the individual test cases for the static and dynamic cases and discusses how these test cases demonstrate the reduction in an observer's ability to predict an agent's behavior.

Personality Framework

The results of the five static test cases for the personality/emotion framework are summarized in Table 18 below. This summary shows, generally, that Extravert/Stability personality dimensions have a slower rate of arousal increase than does the Introvert/Neurotic dimension. This is consistent with expectations, given our implementation. Not demonstrated in these particular test cases is the effect of the Explorer/Preserver dimension on arousal. According to our implementation, the Preserver dimension would experience a faster arousal rate increase than would the Explorer dimension. The reason this is not evidenced in these particular test cases is because the pain/pleasure inputs were so great, that the effects of these inputs outweighed the effects of the clarity/confusion inputs. That is, as indicated in equation 2, the pain/pleasure inputs were always the maximum value and thus made the clarity/confusion inputs insignificant. This is because this set of test cases was static and thus, the opportunity to develop "expectations" for the agent was very short. In the following set of test cases, the dynamic implementation of these cases (1 – 5) does trigger a difference in arousal according to this personality dimension. This is because the cumulative effects of time allow the agent to form expectations that can then be validated or invalidated.

Table 18

Next Arousal Value Summarized Over 5 Static Test Cases

			Next Arousal				
			Case1	**Case 2**	**Case 3**	**Case 4**	**Case 5**
Extraversion	Neuroticism	Explorer	0.51347999	0.53219249	0.56337999	0.56337999	0.58832999
Extraversion	Neuroticism	Preserver	0.51347999	0.53219249	0.56337999	0.56337999	0.58832999
Extraversion	Stability	Explorer	0.51347999	0.51971749	0.53842999	0.53842999	0.54674665
Extraversion	Stability	Preserver	0.51347999	0.51971749	0.53842999	0.53842999	0.54674665
Introversion	Neuroticism	Explorer	0.53842999	0.59456749	0.68812999	0.68812999	0.76297999
Introversion	Neuroticism	Preserver	0.53842999	0.59456749	0.68812999	0.68812999	0.76297999
Introversion	Stability	Explorer	0.53842999	0.55714249	0.61327999	0.61327999	0.63822999
Introversion	Stability	Preserver	0.53842999	0.55714249	0.61327999	0.61327999	0.63822999

Model Output

Also of interest in the static cases is the fact that "Next Arousal" values for Cases 3 and 4 are identical. Further examination reveals that this is because of parameters used to define values of pleasure/pain factors. That is, both Case 3 and Case 4 have maximum Pain value of .6, representing pain experienced for factor "People Shooting at Me". Case 4, where the SOF's teammate is hit, adds the additional Pain factor of "Teammate Hit". Since the pain value assigned to this factor is .5, it is outweighed by the initial value of .6 representing the fact the SOF is being shot at.

The results of the dynamic test case for the personality/emotion framework were presented in Table 17. As discussed in the previous section these cases allow the clarity and confusion signals to influence arousal. Also, these cases demonstrate that the agent starts to consider different options (highlighted in red italics) and actually selects different behaviors based on its personality make-up and state of arousal. Thus, in third event, "Enemy Sees Me", it is apparent that the two agents (ESE and INP) are beginning to consider different options, though they still choose the same option. Then, as the scenario progresses to the fourth event, "Partner-Shot", the action selected by the two agents actually diverges. That is, Agent (ESE) selects to "report-man-injured" and Agent (INP) selected "retreat-to-cover". Of particular interest at this event is the fact that, in the static case, these two cases (tested individually and separately) yielded the same "Next Arousal" values. However, in the dynamic cases, where confusion/clarity parameters are allowed to influence Agents' states, the values are different. This speaks to the power of the method of allowing emotions to emerge. That is, the same set of scenario state variables will not result in the same action by an Agent, but that Agent's action is dependent on history and past experiences in the scenario.

Reducing Agent Predictability

To measure the reduction in predictability of the agent's behavior, we compared the range of the agent's response space using a classic, deterministic state-transition approach with the range of the agent's response space using our emotional model, which is also deterministic. Thus in a classical state-transition construct based on change in world state, as seen in equation 7, there is some fixed number of outputs, given a current world state and an input.

$$\lambda_i = \lambda(q, E_i) \tag{7}$$

where λ_i = set of outputs, for any external input, E_i and any state, q

On the other hand, our approach, still viewed from perspective of state-transition construct, also selects output as a function of the world state and an input. In this case, however, that input is augmented by another state variable internal to the agent (e.g., arousal).

$$\lambda_j = \lambda(q, E_j, I_j) \tag{8}$$

where I_j = is an input internally generated by Agent

To compare the two methods, we manually calculated the number of behaviors possible for each state in the prototype system for the classic approach and the emotions approach. Comparison of these numbers reveals an increased size in response space by average of 3.1, as shown in equation 9.

$$3.1(\lambda_i) \cong \lambda_j \tag{9}$$

What makes this approach useful for generating less predictable behavior is the fact that the additional input is internal to the Agent and thus, not detectable by humans interacting with the scenario. So, for example, while it might be easy for a human participant to learn that "when X happens in the world, the agent will do Z", it is more difficult to learn that "when X happens in the world and the agent's emotional state is Y, the agent will do Z". Primarily, the reason this is difficult to predict is because "Y", the agent's emotional state, is not obvious to the human participant.

Conclusions and Future Work

From a research perspective, future work includes continued experimentation to assist in understanding and improving the theory underlying this model. For example, because of lack of real-world data, the values of factors influencing Arousal, Pain, Pleasure, and Clarity/Confusion were derived through expert opinion of cognitive theorist specializing in emotions research. If formal methods for measuring these values in naturalistic settings were available, they would provide for more confidence in the accuracy of the model.

Future evaluation could explore the effects of emotional situations on agents with various types of background knowledge. In terms of the Army, this might be a way of differentiating the behavior of Soldiers of different ranks. For example, Soldiers with a great deal of experience will not lose access to most of their knowledge even when highly aroused, because their knowledge will be more deeply ingrained. Because arousal serves as a filter on the retrieval and applicability of long-term knowledge, we can certainly expect the same situation to affect differently agents that have individual differences in their long-term knowledge storage. It will be difficult to do a systematic study along these lines, but we can attempt to test agents that include "typical" knowledge differences that might arise from differences in training and background experience.

Ultimately, in the training community, the worth of these models must be measured in terms of improved training. That is, a student interacting with CGFs whose behavior is moderated by emotions should have better performance scores on some training task than a student interacting with "vanilla" CGFs. However, we know of no studies that have investigated whether military training is indeed improved by the use of CGFs with these capabilities. A less ambitious test for model performance has been to determine whether behavior moderators make CGFs appear more realistic (i.e., more humanlike). Of course, since field data to adequately perform such tests do not exist, this becomes a highly subjective measure. Moreover, adopting this measure of performance tacitly and baselessly advances the assumption that CGFs with "more humanlike" behavior (e.g., emotions) will improve training. Again, there is no evidence to support this assumption.

Based on our research, we believe that the incorporation of an emotions model into CGFs can make the behavior of the CGFs less predictable. But, we have not demonstrated, in a pure sense, that emotional CGF behavior is more realistic, nor have we demonstrated that the use of emotional CGFs will improve training. It is our opinion and recommendation that somewhere along this vein of research, funding agencies and sponsors of behavior moderator research formally investigate the assumed benefits of incorporating these models into CGFs.

References

Abell, M. (2001). Soldiers as Distance Learners: What Army Trainers Need to Know. The Proceedings of 2000 Interservice/Industry Training and Education Conference.

Amit, D.J. (1995). The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences,* 18(4): 617-657.

Barnden, (1995). Artificial Intelligence and Neural Networks. In The Handbook of Brain Theory and Neural Networks, Michael A. Arbib, editor. The MIT Press, Cambridge, MA., pp. 98-102.

Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology,* MIT Press, MIT.

Carver, C.S., Sutton, S.K. & Scheier M.F. (2000). Action, emotion and personality: Emerging conceptual integration. *Personality and Social Psychology Bulletin,* Vol 26(6). p. 741-751.

Chapman, C.R. (1995). The affective dimension of pain: A model. *Advances in Pain Research & Therapy,* **22**, p. 283-301.

Chown, E. (1993). The Adaptive Power of Affect: Learning in the SESAME Architecture. in J.A. Meyer, H.L. Roitblat, and S. Wilson (eds.), From Animals to Animats 2: Proceedings of the Second International Conference on Simulation Adaptive Behavior. MIT Press, Cambridge, MA.

Chown, E. (1994). Consolidation and Learning: A Connectionist Model of Human Credit Assignment. Doctoral Dissertation, The University of Michigan.

Chown, E. (2002). Reminiscence and arousal: A connectionist model. Submitted to the Twenty Fourth Annual Meeting of the Cognitive Science Society

Cisneros, J.E., Rajput, S., Karr, C.R., and Parsons, R. (1996). Machine Learning in ModSAF. Report IST-CR-96-12 prepared for STRICOM, Contract No. N61339-92-C-0045.

Costa, P.T., and McCrae, R. R., (1980). Influence of Extraversion and Neuroticism on Subjective Well-Being: Happy and Unhappy People. Journal of Personality and Social Psychology, 38, 6688-678.

Costa, P.T., and McCrae, R. R., (1995). Primary Traits of Eysenck's PEN system: three- and five- factor solutions. *Journal of Personality and Social Psychology,* 69, 308-317.

Craig, K.D. (1999). Emotions and psychobiology. In Wall, P.D. & Melzack, R. (eds.) *Textbook of Pain.* Fourth Edition. Churchill Livingstone, p. 331-343.

Damasio, A. R. (1995). *Descartes' error: Emotion, reason, and the human brain.* New York: Avon.

Davidson, R.J. (1990). Approach-withdrawal and cerebral asymmetry. *Journal of Personality and Social Psychology,* 58, p. 330-341.

Digman, J. M., (1989). Five robust trait dimensions: Development, stability, and utility. Journal of Personality, vol. 57, no. 2, p. 195-214.

Digman, J. M., (1990). Personality structure: Emergence of the five-factor model. *Annual Review of Psychology 1990,* vol. 41, p. 417-440.

D'Ydewalle, G., Ferson, R. & Swerts, A. (1985). Expectancy, Arousal, and Individual Differences in Free Recall. *Journal of Memory and Language*, 519-525.

Epstein, J.M. and Axtell, R. (1996). Growing Artificial Societies: Social Science From the Bottom Up, MIT Press

Eysenck, H.J., and Eysenck, M. (1985). *Personality and Individual Differences: A Natural Science Approach.* New York: Plenum Press.

Eysenck, H.J. (1981). A Model for Personality. New York: Springer.

Eysenck, H.J. (1991). Dimensions of Personality: 16, 5, or 3? Criteria for a Taxonomic Paradigm. Personality and Individual Differences, vol. 12, no., 8, p. 773-790.

Ferguson, E.D. (1982) Motivation: An Experimental Approach. Krueger Publishing Company. Malabar, FL.

Forbell, E., and Chown, E. (2000). Lexical contact during speech perception: A connectionist model. In *Proceedings of the Twenty Second Annual Meeting of the Cognitive Science Society.*

Franceschini, R., McBride, D., and Sheldon, E. (2001). Modeling the Vincennes Incident Using Affective Computer Generated Forces. In Proceedings of the Tenth Conference on Computer Generated Forces. May 15-17, Norfolk, VA., p. 65 - 75.

French, J., Wetteland, C., Bower, S., and Edgin, T. (2001). Estimating Crew Fatigue in Discrete Event Simulations. Proceedings of the Summer Conference on Computer Simulations. Orlando, FL.

Frasconi, P., Gori, M., Maggini, M., and Soda, G. (1995). Unified Integration of Explicit Knowledge and Learning by Example in Recurrent Networks. IEEE Transactions on Knowledge and Data Engineering, vol. 7, no. 2, pp. 340-246.

Gadanho, S. & Hallam, J. (1998). Emotion-triggered learning for autonomous robots. *In Simulation of Adaptive Behavior 1998 Workshop on Grounding Emotions in Intelligent Systems.*

Gratch, J. (2000). Modeling the Interplay Between Emotions and Decision-Making. Proceedings of the 9th Computer Generated Forces and Behavior Representation Conference, Orlando, FL., pp. 399-409.

Gratch, J. (1999). Why You Should Buy and Emotional Planner. In Proceedings of the Agents '99 Workshop on Emotion-based Agent Architectures (EBAA'99) and ISI Research Report ISI/RR-99-465.

Gratch, J., and Marsella, S. (2001). Modeling Emotions in the Mission Rehearsal Exercise. Proceedings of the 10th Computer Generated Force and Behavior Representation Conference, May, Norfolk, VA. Pp. 457-465.

Gratch, J. (2000). Emile: Marshalling Passions in Training and Education. Proceedings of the 4th International Conference on Autonomous Agents, Barcelona, Spain.

Gross, J.J., Sutton, S.K. & Ketelaar, T. (1998). Relations between affect and personality: Support for the affect-level and affective reactivity views. *Personality and Social Psychology Bulletin,* Vol 24(3), p. 279-288.

Hebb, D.O. (1972). Emotional Arousal and Performance. In *Textbook of Psychology,* 3rd ed. Philadelphia: W.B. Saunders Co.

Hebb, D.O. (1949). The organization of behavior. New York: Wiley.

Heeringa, B. and Cohen, P. (2000). An Underlying Model for Defeat Mechanisms. Proceedings of the 2000 Winter Simulation Conference. Dec., Orlando, FL., pp. 933-939.

Henninger A., Gonzalez, A., Gerber, W., , Georgiopoulos, M., and DeMara, R., (2000). On the Fidelity of SAFs: Can Performance Data Help, Proceedings of the '00 Interservice/Industry Training, Simulation and Education Conference (I/ITSEC). Orlando, FL., November, 2000.

Hill, R. W., Chen, J., Gratch, J., Rosenbloom, P., & Tambe, M. (1997). Intelligent agents for the synthetic battlefield: A company of rotary-wing aircraft. In *Proceedings of the Ninth Conference on Innovative Applications of Artificial Intelligence,* 1006-1012. Menlo Park, CA: AAAI Press.

Hudlicka, E., and Billingsley, J. (1999). Representing Behavior Moderators in Military Human Performance Models. Proceedings of 8th Conference on Computer Generated Forces and Behavior Representations, Orlando, FL.

Ilachinski, A. (1996). Land Warfare and Complexity, Part I: Mathematical Background and Technical Sourcebook, Center for Naval Analyses Information Manual CIM-461, July 1996, Unclassified.

Ilachinski, A. (1996). Land Warfare and Complexity, Part II: An Assessment of the Applicability of Nonlinear Dynamics and Complex Systems Theory to the Study of Land Warfare, Center for Naval Analyses Research Memorandum CRM-68, July 1996, Unclassified.

James, W. (1892).*Psychology: The briefer course* (1961 ed.). New York: Harper & Row.

Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., & Koss, F. V. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine, 20*(1), 27-41.

Jones, R.M., Laird, J.E., and Neville, K. (1997). Modeling pilot fatigue with a synthetic behavior model. Proceedings of the 7th Conference on Computer Generated Forces and Behavioral Representation, Orlando, FL. P. 349-357.

Kaplan, S., Weaver, M., & French, R. (1990). Active symbols and internal models: Towards a cognitive connectionism. *AI & Society,* **4**, p. 51-71.

Kaplan, S., & Kaplan, R. (1982). *Cognition and Environment: Functioning in an Uncertain World.* Ulrichs Books: Ann Arbor, MI, 1989 edition.

Kaplan, S. (1991). Beyond rationality: Clarity-based decision making. In T. Garling and G. Evans (Eds.), *Environment, Cognition, and Action: An Integrative Multidisciplinary Approach,* pp. 171-190. New York: Oxford University Press.

Kaplan, S., Sonntag, M. & Chown, E. (1991) Tracing Recurring Activity in Cognitive Elements (TRACE): A Model of Temporal Dynamics in a Cell Assembly. Connection Science 3, 179-206.

Keirsey, D.W. and Bates, M. (1984). Please Understand Me: Character and Temperament Types. Matrix Books, Inc.

Lazarus, (1984). On the primacy of cognition. American Psychologist, 39, 124-129.

LeDoux, J.E. (1992). Brain Mechanisms of Emotion and Emotional Learning. *Current Opinions in Neurobiology, 2,* 191-197.

LeDoux, J. E., and Fellous, J. (1995). Emotion and Computational Neuroscience. In The Handbook of Brain Theory and Neural Networks, Michael A. Arbib, editor. The MIT Press, Cambridge, MA., pp. 356 – 360.

Levine, D.S., and Leven, S. J., (1992). Motivation, Emotion, and Goal Direction in Neural Networks, Erlbaum, Hillsdale, NJ.

Lindsay, P.H., and Norman, D.A. (1977). Human Information Processing. Academic Press, Inc. New York, NY.

McCrae, R.R., and Costa, P.T. (1989). Reinterpreting the Myers-Briggs Type Indicator from the perspective of the five-factor model of personality. Personality and Individual Differences, 6, 587-597.

McKenzie, F., Catanzaro, J., and Petty, M.D. (2001). A Personality-Based Command Decision-Maker: Results and Recommendations. In Proceedings of the Tenth Conference on Computer Generated Forces. May 15-17, Norfolk, VA., p., 351-359.

Merskey, H. (1979). Pain terms: A list with definitions and a note on usage. Recommended by the International Association for the Study of Pain (IASP) Subcommittee on Taxonomy. *Pain,* **6,** pp 249-252.

Meyer, G.J., and Shack, J.R., (1989). Structural Convergence of Mood and Personality: Evidence for Old and New Directions. Journal of Personality and Social Psychology, vol. 57, p. 491-506.

Milner, P.M. (1991). Brain Stimulation Reward: A Review. *Canadian Journal of Psychology,* vol. 45, p. 1-36.

Myers, I.B., and McCaulley, M.H. (1985). Manual: A guide to the development and use of the Myers-Briggs Type Indicator. Palo Alto, CA: Consulting Psychology Press.

Newell, A. 1990. *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Nielsen, P., Smoot, D., Martinez, R., and Dennison, J. (2000). Participation of TacAir-Soar in Road Runner and Coyote Exercises at Air Force Research Lab, Mesa, AZ. Proceedings of the 9th Conference on Computer Generated Forces and Behavioral Representation, Orlando, FL., p. 173-180.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, NJ: Prentice Hall.

Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions.* Cambridge, England: Cambridge University Press.

Pew, R.W., and Mavor, A.S., eds. (1998). Modeling Human and Organizational Behavior: Application to Military Simulations. Washington, DC: National Academy Press.

Ploghaus, A., Tracey, I., Gati, J.S., Clare, S., Menon, R.S., Mathews, P.M., & Rawlins, N.P. (1999). Dissociating pain from its anticipation in the human brain. *Science,* **284,** 1979-1985.

Revelle, W. & Loftus, D. (1990) Individual differences and arousal: implications for the study of mood and memory. *Cognition and Emotion*, 4, 209-237.

Robinson, R.G., & Coyle, J.T. (1980). The differential effect of right vs. left hemisphere cerebral infarction on catecholamines and behavior in the rat. *Brain Research*, 188, pp. 63-78

Ropella, G.E.P., Railsback, S.F., and Jackson, S. K. (2002). Software Engineering Considerations for Individual-based Models. *Natural Resource Modeling.* Vol(15), no(1).

Rykiel, E. (1995). Testing ecological models: the meaning of validation. *Ecological Modeling. Vol(*90), pp. 229 – 244.

Rosenberg, A. & Kagan, J. (1987). Iris pigmentation and behavioral inhibition. *Developmental Psychobiology*, 20, p. 377-392.

Rosenbloom, P., Laird, J., and Newell, A (1993). The Soar Papers: Research on Integrated Intelligence. MIT Press, Cambridge, MA. 1993.

Rusting, CL, and Larsen, R.J., (1996). Extraversion, Neuroticism, and Susceptibility to Positive and Negative Affect: A Test of Two Theoretical Models. Journal of Personality and Individual Differences., vol., 22, No. 5, p 607-612.

Saggino, A., and Kline, P. (1996). The location of the Myers-Briggs Type Indicator in personality factor space. Personality and Individual Differences, 21, 591-597.

Salovey, P., & Mayer, D. (1990). Emotional intelligence. *Imagination, Cognition and Personality*, **9**(3): pp. 185-211.

Schachter, S. and Singer, J. E. (1962). Cognitive, social and physiological determinants of emotional state. *Psychological Review*, 69: 379-399.

Schultz, D. (1981). A History of Modern Psychology. Academic Press, New York, NY.

Schwarz, N., and Clore, G.L. (1983). Mood, misattribution, and judgments of well-being; Information and directive function of affective states. Journal of Personality and Social Psychology, 45, 513-523.

Shavlik, J.W., Mooney, R.J., and Towell, G.G. (1991). Symbolic and Neural Learning Algorithms: An Experimental Comparison. Machine Learning, vol. 6, no. 2, pages 111-143.

Smith, S., and Petty, M. (1992). Controlling Autonomous Behavior in Real-Time Simulation. In Proceedings of the Second Conference in Computer Generated Forces and Behavior Representation. Orlando, FL., May, 1992.

Sonntag, M.. (1991). Learning sequences in an associative network: A step towards cognitive structure. Doctoral Dissertation. The University of Michigan.

Sun, R., and Alexandre, F. (1997). Connectionist-Symbolic Integration: From Unified to Hybrid Approaches, Lawrence Erlbaum Associates, Mahwah, N.J.

Tan, A. H. (1997). Cascade ARTMAP: Integrating Neural Computation and Symbolic Knowledge Processing. IEEE Transactions on Neural Networks, vol. 8, no. 2, pages 237-250.

Taylor, G., Koss, F., Nielsen, P. (2001). Special Operations Forces IFORs. I*n Proceedings of the Tenth Conference on Computer Generated Forces and Behavioral Representation.* May 15-17, Norfolk, VA. pp. 301 – 306.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185, p. 1124-1131.

Velasquez, J.D. (1997). Modeling emotions and other motivations in synthetic agents. In: *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97)*. Providence, RI: MIT/AAAI Press.

Velasquez, J.D. (1998). When robots weep: Emotional memories and decision-making. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*. Madison, WI: MIT/AAAI Press.

Watson, D., & Tellegen,, A. (1985). Towards a consensual structure of mood. *Psychological Bulletin*, 98, p. 219-235.

Williams, D.G. (1989). Neuroticism and Extraversion in Different Factors of the Affect Intensity Measure. Personality and Individual Differences, vol. 10, p. 1095-1100.

Appendix A. Emotion Model Expressed in Soar Rules

```
sp {top-ps*elaborate*state*emotions
  (state <s> ^problem-space.name top-ps)
-->
  (<s> ^emotions <em>)
}

sp {top-ps*elaborate*emotions*arousal*baseline
  (state <s> ^emotions.arousal <a>)
-->
  (<a> ^baseline 0.5)
}

sp {top-ps*elaborate*emotions*recovery-rate
  (state <s> ^emotions.arousal <a>)
  (<a> ^magnitude <am>
     ^baseline <bl>)
-->
  (<a> ^recovery-rate (- 0.501 (abs (- <am> <bl>))))
}

sp {top-ps*elaborate*emotions*arousal*new-magnitude
  (state <s> ^emotions. <em>)
  (<em> ^arousal <a>
      ^pleasure <pl>
      ^pain <pa>)
  (<a> ^recovery-rate <rr>
     ^baseline <bl>
     ^magnitude <am>
     ^stimulus-magnitude <asm>
     ^sensitivity <as>)
  (<pl> ^delta-magnitude <pldm>)
  (<pa> ^delta-magnitude <padm>)
-->
  (<a> ^new-magnitude (+ (* <rr> (- <bl> <am>))
                (* <as> (+ <pldm> <padm> <asm>))))
}

sp {top-ps*elaborate*emotions*pleasure-pain*delta-magnitude
  (state <s> ^emotions.<< pleasure pain >> <pl>)
  (<pl> ^magnitude <plm>
      ^last-magnitude <pllm>)
-->
  (<pl ^delta-magnitude (- (<plm> <pllm>)))
}

sp {top-ps*elaborate*emotions*pleasure*stimulus*magnitude
  (state <s> ^emotions.<< pleasure pain clarity confusion arousal >> <pl>)
  (<pl> ^sensitivity <pls>
      ^stimulus <stim>)
  (<stim> ^signal <signal>)
-->
  (<stim> ^magnitude (* <pls> <signal>))
}
```

```
sp {top-ps*elaborate*emotions*pleasure*stimulus-magnitude
  (state <s> ^emotions.<< pleasure pain clarity arousal >> <pl>)
-->
  (<pl> ^stimulus-magnitude ;## some function of stimulus magnitudes)
}

sp {top-ps*elaborate*emotions*pleasure*new-magnitude
  (state <s> ^emotions <em>)
  (<em> ^pleasure <pl>
      ^pain.stimulus-magnitude <pa>)
  (<pl> ^sensitivity <pls>
      ^stimulus-magnitude <plsm>)
-->
  (<pl> ^new-magnitude (* <pls> (max (- <plsm> <pasm>) 0)))
}

sp {top-ps*elaborate*emotions*pain*new-magnitude
  (state <s> ^emotions <em>)
  (<em> ^pain <pa>
      ^pleasure.stimulus-magnitude <plsm>)
  (<pa> ^sensitivity <pas>
      ^stimulus-magnitude <pasm>)
-->
  (<pa> ^new-magnitude (* <pas> (max (- <pasm> <plsm>) 0)))
}

sp {top-ps*elaborate*emotions*clarity*new-magnitude*clear
  (state <s> ^emotions <em>)
  (<em> ^clarity <cl>)
  (<cl> ^clarity-sensitivity <cls>
      ^stimulus-magnitude { <clsm> >= 0 })
-->
  (<cl> ^new-magnitude (* <cls> <clsm>))
}

sp {top-ps*elaborate*emotions*clarity*new-magnitude*confused
  (state <s> ^emotions.clarity <cl>)
  (<cl> ^confusion-sensitivity <cls>
      ^stimulus-magnitude { <clsm> < 0 })
-->
  (<cl> ^new-magnitude (* <cls> <clsm>))
}

sp {top-ps*elaborate*emotions*pleasure*stimulus*clarity
  (state <s> ^emotions <em>)
  (<em> ^clarity <cl>
      ^pleasure <pl>)
  (<cl> ^magnitude { <clm> > 0})
-->
  (<pl> ^stimulus <pls> + &)
  (<pls> ^signal <clm>
      ^object *clarity*)
}

sp {top-ps*elaborate*emotions*pain*stimulus*clarity
```

```
  (state <s> ^emotions <em>)
  (<em> ^clarity <cl>
       ^pain <pa>)
  (<cl> ^magnitude { <clm> < 0})
-->
  (<pa> ^stimulus <pas> + &)
  (<pas> ^signal (abs <clm>)
        ^object *confusion*)
}
```

Appendix B. SOF-Soar Rules Added for Cognitive Reactions

```
### Consider shoot if I think an enemy sees me
sp {any-ps*propose*shoot
  (state <s> ^top-state <ts>)
  (<ts> ^stimulus <s1> <s2> <s3>)
  (<s1> ^agent-id <x>
      ^type sees-me
      ^value *yes*)
  (<s2> ^agent-id <x>
      ^type exists
      ^value *yes*)
  (<s2> ^agent-id <x>
      ^type is-enemy
      ^value *yes*)
-->
  (<s> ^operator <o> +, =)
  (<o> ^name shoot
     ^at-agent <x>
     ^new-terminate *yes*)}


### Consider run-away if I think an enemy sees me (and I'm not already
### running)
sp {any-ps*propose*run-away
  (state <s> ^top-state <ts>)
  (<ts> ^stimulus <s1> <s2> <s3>)
  (<s1> ^agent-id <x>
      ^type sees-me
      ^value *yes*)
  (<s2> ^agent-id <x>
      ^type exists
      ^value *yes*)
  (<s2> ^agent-id <x>
      ^type is-enemy
      ^value *yes*)
 -{(<ts> ^stimulus <sr>)
  (<sr> ^type running
      ^value *yes*)
}
-->
  (<s> ^operator <o> +, =)
  (<o> ^name run-away
     ^from-agent <x>
     ^new-terminate *yes*)
}

### Consider look-toward-shot if I heard a shot, and I don't already
### know where shots are coming from
sp {any-ps*propose*look-toward-shot
  (state <s> ^top-state <ts>)
  (<ts> ^stimulus <s1> <s2> <s3>)
  (<s1> ^type shot-heard
      ^value *yes*
```

```
        ^direction <dir>
       -^attended *yes*)
   (<s2> ^agent-id <x>
        ^type source-of-shots
        ^value <> *unknown*)
   (<s3> ^agent-id <x>
        ^type location
        ^value *unknown*)
-->
   (<s> ^operator <o> +, =)
   (<o> ^name look-toward-shot
       ^direction <dir>
       ^new-terminate *yes*)
}

### Consider report-im-alive if my leader does not know I'm alive
sp {any-ps*propose*report-im-alive
   (state <s> ^top-state <ts>)
   (<ts> ^stimulus <s1>
        ^leader-id <x>)
   (<s1> ^agent-id <x>
        ^type knows-im-alive
        ^value *no*)
-->
   (<s> ^operator <o> +, =)
   (<o> ^name report-Im-alive
       ^to-agent <x>
       ^new-terminate *yes*)
}

### Consider report-enemy-contact if an enemy exists and I haven't reported
### it
sp {any-ps*propose*report-enemy-contact
   (state <s> ^top-state <ts>)
   (<ts> ^stimulus <s1> <s2>)
   (<s1> ^agent-id <x>
        ^type exists
        ^value *yes*)
   (<s2> ^agent-id <x>
        ^type is-enemy
        ^value *yes*)
 -{(<ts> ^stimulus <sr>)
   (<sr> ^agent-id <x>
        ^type reported-contact
        ^value *yes*)
}
-->
   (<s> ^operator <o> +, =)
   (<o> ^name report-enemy-contact
       ^to-agent <x>
       ^new-terminate *yes*)
}

### freeze is a default action that always gets proposed, but will likely
### get filtered out, unless arousal is very high.  Even when it's not
### filtered, other actions will be preferred to it (unless they get
```

```
### filtered out)
sp {any-ps*propose*freeze
  (state <s> ^type state)
-->
  (<s> ^operator <o> +)
  (<o> ^name freeze)
}
```

Appendix C. Interpretation of Rules for SOF Cognitive Reactions

1. If **(enemy-exists = *yes*)** AND **(leader-knows-observation = *no*)**, then establish immediate effect of **(leader-knows-observation = *unknown*)** and establish expectation of **(leader-knows-observation = *yes*)** and propose transition to "**report-observation**".

If Agent sees enemy and Agent hasn't already filed an observation report on that enemy, then Agent will proceed to report the observation. Initially, the Agent does not know if the report has been transmitted successfully, but the Agent expects that the report will be transmitted and the Agent expects that the Commander will be informed of this observation.

2. If **(enemy-exists = *yes*)** AND **(enemy-sees-me = yes)** AND **(leader-knows-enemy-contact = *no*)** then establish immediate effect of **(leader-knows-enemy-contact = *unknown*)** and establish expectation of **(leader-knows-enemy-contact = *yes*)** and propose transition to "**report-enemy-contact**".

If Agent sees enemy and Agent has determined that the enemy has seen Agent and the Agent has not already reported this contact with the enemy, then Agent will proceed to report the enemy contact. Initially, the Agent does not know if the report has been transmitted successfully, but the Agent expects that the report will be transmitted and the Agent expects that the Commander will be informed of this contact.

3. If **(leader-knows-im-alive = *no*)**, then establish immediate effect of **(leader-knows-im-alive = *unknown*)** and establish expectation of **(leader-knows-im-alive = *yes*)** and propose transition to "**report-im-alive**".

If Agent_{sub} determines that Agent_{sup} does not know whether or not the Agent_{sub} is alive, then Agent_{sub} will proceed to report that he "is alive". Initially, Agent_{sub} does not know if the report has been transmitted successfully, but Agent_{sub} expects that the report will be transmitted successfully and Agent_{sub} expects that Agent_{sup} will know that Agent_{sub} is alive.

4. If **(leader-knows-im-alive = *no*)**, then establish immediate effect of **(leader-knows-im-alive = *unknown*)** and **(enemy-sees-me = *yes*)** and establish expectation of "**leader-knows-im-alive = *yes*)** and propose transition to "**shout**".

If Agent_{sub} determines that the Agent_{sup} does not know whether or not the Agent_{sub} is alive, then Agent_{sub} will proceed to "shout". Initially, the Agent_{sub} does not know if the shout has been heard by Agent_{sup}, but Agent_{sub} expects that the shout will be heard by Agent_{sup} and Agent_{sub} expects that the shout will also be heard by the enemy.

5. If **(enemy-exists = *yes*)** AND **(enemy-sees-me = *yes*)**, then establish immediate effect of **(enemy-exists = *unknown*)** and **(enemy-sees-me = *unknown*)** and establish expectation of **(enemy-exists = *no*)** and **(enemy-sees-me = *no*)** and propose transition to "**shoot**".

If Agent sees enemy and Agent has determined that the enemy has seen Agent and the Agent has not already reported this contact with the enemy, then Agent will proceed to shoot Initially, the Agent does not know if the shot will be successful, but the Agent expects that the shot will be successful and that it will damage the enemy such that the enemy threat no longer exists and the enemy can no longer see the Agent.

6. If **(heard-shot = *yes*)** AND **(enemy-location = *unknown*)**, then establish expectation of **(enemy-location = shot-location(X,Y))** and propose transition to "**search-for-shooter**".

If Agent hears gun shot and the Agent does not know where the enemy is located, then Agent will proceed to search for the shooter The Agent expects that the location of where shot was originated will yield the location of the enemy.

7. If **(enemy-exists = *yes*)** AND **(enemy-sees-me = *yes*)**, then establish immediate effect of **(enemy-sees-me = *unknown*)** and **(leader-knows-im-alive = *no*)** and establish expectation of **(enemy-sees-me = *no*)** and propose transition to **"retreat-to-cover"**.

If $Agent_{sub}$ sees enemy and $Agent_{sub}$ has determined that the enemy has seen $Agent_{sub}$, then $Agent_{sub}$ will proceed to retreat. Initially, $Agent_{sub}$ does not know if the retreat has provided cover/concealment. Also, initially, $Agent_{sub}$ recognizes that $Agent_{sup}$ may not know that $Agent_{sub}$ is alive. But the $Agent_{sub}$ expects that the retreat will result in the enemy no longer being able to see $Agent_{sub}$.

8. If **"perceive-partner-bullet-pain"** AND **(leader-knows-man-injured = no),** then establish immediate effect of **(leader-knows-man-injured = *unknown*)** and establish expectation of **(leader-knows-man-injured = *yes***), and propose transition to **"report-man-injured"**.

If Agent1 perceives bullet pain in Agent 2, then Agent1 will proceed to report that his partner is injured. Initially, Agent1 does not know if that report has been successfully transmitted. But, Agent1 expects that the report will be transmitted successfully and Agent1 expects that the Commander will be informed of Agent1's injury.

<u>Observation Team: Enumerated Cognitive Reactions</u>

- **report-observation**
- **report-man-injured**
- **report-enemy-contact**
- **report-I'm-alive**
- **shoot**
- **search-for-shooter**
- **retreat-to-cover**
- **shout**

Appendix D. Scenario Assumptions

Of total number of state variable combinations, certain combinations of these variables cannot exist within context of the scenario. These combinations are listed below as model assumptions.

If **(enemy-exists** = ***no***) AND **(leader-knows-observation** = ***yes*** OR **leader-knows-observation** = ***unknown***)

This assumption controls cases where the leader knows about an observation, but an enemy doesn't exist.

If **(enemy-sees-me** = ***no*** OR **enemy-sees-me** = ***unknown***) AND **(leader-knows-enemy-contact** = ***yes*** OR **leader-knows-enemy-contact** = ***unknown***)

This assumption controls cases where Agent hasn't been spotted by enemy, but leader knows there was contact.

If **(perceive-partner-bullet-pain** = ***no***) AND **(leader-knows-man-injured** = ***yes*** OR **leader-knows-man-injured** = ***unknown***)

This assumption controls cases where Agent's partner hasn't been shot, but leader believes Agent's partner is shot.

If **(heard-shot = no)** AND **(perceive-partner-bullet-pain = yes** OR **(leader-knows-man-injured = yes** OR **leader-knows-man-injured = unknown))**

This assumption controls cases where Agent has not heard a shot, but Agent's partner has bullet pain OR leader believes Agent's partner is injured. So, for example, this would mean that our scenario could not support any weapon but a gun that makes sound.

If **(enemy-exists = no)** AND **(enemy-sees-me = yes)**

This assumption controls cases where the enemy can see me, but no enemy exists.

If **(enemy-exists = no)** AND **(enemy-location = (X,Y))**

This assumption controls cases where an enemy doesn't exist, but Agent knows enemy's location.

If **(leader-knows-enemy-contact = no)** AND **(leader-knows-man-injured = yes** OR **leader-knows-man-injured = unknown)**

This assumption controls cases where a leader doesn't know about enemy contact, but does believe an Agent's partner is injured. In other words, if there was an injury, it cam from enemy-contact.

Appendix E. Rules for Arousal Thresholds

```
####################################
### Here are the general rules for rejecting proposed operators based on
### their arousal thresholds. This is our simulation of "cognitive
### focus of attention". Operators get proposed based on normal
### decision making, then get filtered out if current arousal conditions ### dictate that a particular action would not be considered. This is ### probably not what's really going on in the brain, but it has the
### functional effect of what our theory says: retrieval of potential ### actions is a function both of knowledge-based retrieval and arousal-### based activation/filtering.

sp {reject*arousal-low-threshold
  (state <s> ^operator <o> +
         ^top-state.emotions.arousal.magnitude <= <th>)
  (<o> ^arousal-low-threshold <th>)
-->
  (<s> ^operator <o> -)
}

sp {reject*arousal-high-threshold
  (state <s> ^operator <o> +
         ^top-state.emotions.arousal.magnitude >= <th>)
  (<o> ^arousal-high-threshold <th>)
-->
  (<s> ^operator <o> -)
}


####################################
### Below are the arousal threshold parameters for the key actions in
### our experimental scenario

sp {shoot*elaborate*arousal-low-threshold
  (state <s> ^operator <o> +)
  (<o> ^name shoot)
-->
  (<o> ^arousal-low-threshold 0.7)
}

sp {shoot*elaborate*arousal-high-threshold
  (state <s> ^operator <o> +)
  (<o> ^name shoot)
-->
  (<o> ^arousal-high-threshold 0.95)
}

sp {run-away*elaborate*arousal-low-threshold
  (state <s> ^operator <o> +)
  (<o> ^name run-away)
-->
  (<o> ^arousal-low-threshold 0.75)
}
```

```
sp {run-away*elaborate*arousal-high-threshold
  (state <s> ^operator <o> +)
  (<o> ^name run-away)
-->
  (<o> ^arousal-high-threshold 0.99)
}

sp {freeze*elaborate*arousal-low-threshold
  (state <s> ^operator <o> +)
  (<o> ^name freeze)
-->
  (<o> ^arousal-low-threshold 0.95)
}

### freeze has no high threshold

### look-towards-shot has no low threshold

sp {look-towards-shot*elaborate*arousal-high-threshold
  (state <s> ^operator <o> +)
  (<o> ^name look-towards-shot)
-->
  (<o> ^arousal-high-threshold 0.8)
}

### "Normal" operators with default arousal values:
###   report-I'm-alive, report-enemy-contact, report-observation
###   report-man-injured

sp {normal-arousal*elaborate*arousal-low-threshold
  (state <s> ^operator <o> +)
  (<o> ^normal-arousal *yes*)
-->
  (<o> ^arousal-low-threshold 0.6)
}

sp {normal-arousal*elaborate*arousal-high-threshold
  (state <s> ^operator <o> +)
  (<o> ^normal-arousal *yes*)
-->
  (<o> ^arousal-high-threshold 0.8)
}

sp {report-I'm-alive*elaborate*normal-arousal
  (state <s> ^operator <o> +)
  (<o> ^name report-I'm-alive)
-->
  (<o> ^normal-arousal *yes*)
}

sp {report-enemy-contact*elaborate*normal-arousal
  (state <s> ^operator <o> +)
  (<o> ^name report-enemy-contact)
-->
  (<o> ^normal-arousal *yes*)
```

```
}

sp {report-observation*elaborate*normal-arousal
  (state <s> ^operator <o> +)
  (<o> ^name report-observation)
-->
  (<o> ^normal-arousal *yes*)
}

sp {report-man-injured*elaborate*normal-arousal
  (state <s> ^operator <o> +)
  (<o> ^name report-man-injured)
-->
  (<o> ^normal-arousal *yes*)
}
```

Appendix F. Rules for Post-filtering Priority Scheme

```
### These are knowledge-based comparisons of potential actions that
### Have made it through the "focus-of-attention" filter provided by
### arousal thresholds. The focus-of-attention filter rejects some
### potential actions, simulating the effect that those actions will not
### even be considered in particular situations. Whatever potential
### actions remain can be deliberately chosen between using decision-
### making knowledge. That happens here. For current testing purposes
### the decision-making comparisons are somewhat simplified.

##################################
### Here are the general rules for knowledge-based comparison of actions
### that have made it through the arousal filtering mechanism. The
### current implementation is a set of relatively simple Soar preferences
### based on priorities. More sophisticated knowledge-based decisions
### would support much more complex kinds of behavior, but we want to
### control things somewhat for terms of running experiments.


sp {prefer*emotions-priority*highest*best
  (state <s> ^operator <o> +)
  (<o> ^emotions-priority highest)
-->
  (<s> ^operator <o> >)
}

sp {prefer*emotions-priority*high
  (state <s> ^operator <o1> + <o2> +)
  (<o1> ^emotions-priority high)
  (<o2> ^emotions-priority { <> highest <> high })
-->
  (<s> ^operator <o1> > <o2>)
}

sp {prefer*emotions-priority*medium
  (state <s> ^operator <o1> + <o2> +)
  (<o1> ^emotions-priority medium)
  (<o2> ^emotions-priority { <> highest <> high <> medium })
-->
  (<s> ^operator <o1> > <o2>)
}

sp {prefer*emotions-priority*low
  (state <s> ^operator <o1> + <o2> +)
  (<o1> ^emotions-priority low)
  (<o2> ^emotions-priority { <> highest <> high <> medium <> low })
-->
  (<s> ^operator <o1> > <o2>)
}

sp {prefer*emotions-priority*very-low
  (state <s> ^operator <o1> + <o2> +)
  (<o1> ^emotions-priority very-low)
```

```
  (<o2> ^emotions-priority { <> highest <> high <> medium <> low <> very-low })
-->
  (<s> ^operator <o1> > <o2>)
}

sp {prefer*emotions-priority*lowest
  (state <s> ^operator <o> +)
  (<o1> ^emotions-priority lowest)
-->
  (<s> ^operator <o1> <)
}

### Catch-all...if there's anything left over, pick at random.
### (But the way we've arranged things for our experiments, this
### shouldn't happen)
sp {prefer*emotions-priority*any*indifferent
  (state <s> ^operator <o> +)
  (<o> ^emotions-priority <don't-care>)
-->
  (<s> ^operator <o> =)
}

######################################
### Below are the priority parameters attached to the significant
###  actions in our study.
###  "highest" suggests a "pure reflex" type of action
###  "high" suggests relatively important actions that should be taken
###       without much though
###  "medium" suggests actions that involve a "normal" amount of
###       deliberation
###  "low" and "very-low" suggest more deliberate actions that might not
###      be taken in high priority situations
###  "lowest" suggests default actions that will only be executed when
###      the agent can think of nothing better to do

sp {look-toward-shot*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name look-towards-shot)
-->
  (<o> ^emotions-priority highest)
}

sp {report-enemy-contact*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name report-enemy-contact)
-->
  (<o> ^emotions-priority high)
}

sp {report-observation*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name report-observation)
-->
  (<o> ^emotions-priority high)
}
```

```
sp {report-I'm-alive*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name report-I'm-alive)
-->
  (<o> ^emotions-priority medium)
}

sp {report-man-injured*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name report-man-injured)
-->
  (<o> ^emotions-priority medium)
}

sp {shoot*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name shoot)
-->
  (<o> ^emotions-priority low)
}

sp {run-away*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name run-away)
-->
  (<o> ^emotions-priority very-low)
}

sp {freeze*elaborate*priority
  (state <s> ^operator <o> +)
  (<o> ^name freeze)
-->
  (<o> ^emotions-priority lowest)
}
```

Appendix G. Interpretation of Arousal Rules

conditions based purely on emotions, not on cognition	preference ordering:
shoot: arousal > .7 AND arousal < .95 run-away arousal > .75 arousal < .99 freeze arousal > .95 look-towards-shot arousal < .8 report-I'm-alive arousal > .3 AND arousal < .8 report-enemy-contact arousal > .3 AND arousal < .8 report-observation arousal > .3 AND arousal < .8 report-man-injured arousal > .3 AND arousal < .8	Highest: look-towards-shot Retreat_to_Cover High: report-enemy-contact, report-observation Middle: report-I'm-alive, report-man-injured Low: shoot Very low: run-away Lowest: freeze

Emotional Reactions:

9. If **(arousal > .7)** AND **(arousal < .95)** AND "shoot" is selected, then propose **"Shoot"**.
10. If **(arousal > .7)** AND **(arousal < .95)** AND **"Retreat_to_Cover"** is selected, then propose **"Retreat_to_Cover"**.
11. If **(arousal > .75)** AND **(arousal < .99)**, then propose **"flee"**.
12. If **(arousal > .95)**, then propose **"freeze"**.
13. If **(arousal < .8)** AND **Search_for_Shooter** is proposed, then propose **"look-towards-shot"**.
14. If **"look-towards-shot"** is proposed, select **look-towards-shot**.
15. If **(arousal > .3)** AND **(arousal < .8)** AND **"report-observation"** is proposed AND **"look-towards-shot"** is NOT proposed, then **"report-observation"**.
16. If **(arousal > .3)** AND **(arousal < .8)** AND **"report-enemy-contact"** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed, then **"report-enemy-contact"**.
17. If **(arousal > .3)** AND **(arousal < .8)** AND **"report-I'm-alive"** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed AND **report-enemy-contact** is NOT proposed, then **"report-I'm-alive"**.
18. If **(arousal > .3)** AND **(arousal < .8)** AND **"report-man-injured"** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed AND **report-enemy-contact** is NOT proposed AND **report-man-injured** is NOT proposed, then **"report-I'm-alive"**.
19. If **shoot** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed AND **report-enemy-contact** is NOT proposed AND **report-man-injured** is NOT proposed, then **"shoot"**.
20. If **Run-Away** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed AND **report-enemy-contact** is NOT proposed AND **report-man-injured** is NOT proposed AND **shoot** is NOT proposed, then **"Run-Away"**.
21. If **Freeze** is proposed AND **"look-towards-shot"** is NOT proposed AND **report-observation** is NOT proposed AND **report-enemy-contact** is NOT proposed AND **report-man-injured** is NOT proposed AND **shoot** is NOT proposed AND **Run_Away** is NOT proposed, then **"Freeze"**.

List of Emotional Reactions for Observation Team: shoot, run-to-cover, flee, freeze, look-toward-shot, report-observation, report-man-injured, report-enemy-contact, report-I'm-alive.